



Europäisches Patentamt
European Patent Office
Office européen des brevets



⑪ Publication number:

0 150 084 B1

⑫

EUROPEAN PATENT SPECIFICATION

⑬ Date of publication of patent specification: 29.05.91 ⑭ Int. Cl. 5: G06F 13/12, G06F 13/38,

H04L 12/42, G06F 13/362

⑮ Application number: 85300009.9

⑯ Date of filing: 02.01.85

⑭ Architecture for intelligent control of data communication adapters.

⑮ Priority: 03.01.84 US 567595
03.01.84 US 567596
03.01.84 US 567597
03.01.84 US 567598
03.01.84 US 567599
03.01.84 US 567730

⑯ Date of publication of application:
31.07.85 Bulletin 85/31

⑭ Publication of the grant of the patent:
29.05.91 Bulletin 91/22

⑮ Designated Contracting States:
DE FR GB

⑯ References cited:
EP-A- 0 054 077
EP-A- 0 095 363
WO-A-84/03192
US-A- 4 335 426

⑭ Proprietor: TEXAS INSTRUMENTS INCORPORATED
13500 North Central Expressway
Dallas Texas 75265(US)

⑮ Inventor: Patrick, Michael W.
7520 Cook Road No. 1907
Houston Texas 77072(US)
Inventor: Buchanan, George
7 Carters Close Sherrington
Newport Pagnell Bucks, MK16 9NW(GB)
Inventor: Laffitte, David S.
15 Chestnut Street
Wellesley Hills Massachusetts 02181(US)
Inventor: Shore, Brian
18 Rotten Row
Riseley Bedford MK44 1E5(GB)
Inventor: Walters, Donald M. Jr.
6010 Elm, Apt 991
Houston Texas 77081(US)
Inventor: Killebrew, Carrell R. Jr.
3034 Pasture Lane
Sugar Land Texas 77479(US)

EP 0 150 084 B1

Note: Within nine months from the publication of the mention of the grant of the European patent, any person may give notice to the European Patent Office of opposition to the European patent granted. Notice of opposition shall be filed in a written reasoned statement. It shall not be deemed to have been filed until the opposition fee has been paid (Art. 99(1) European patent convention).

Inventor: Fanini, Otto N.
10027 Splice Lane No.510
Houston Texas 77092(US)
Inventor: Chang, Ki Suk
9901 S. Wilcrest Dr.
Houston Texas 77099(US)
Inventor: Sacarsen, Stephen P.
12952 Wrevine
Houston Texas 77072(US)
Inventor: Stambaugh, Mark A.
9901 S. Wilcrest Drive
Houston Texas 77099(US)

⑥ Representative: Abbott, David John et al
Abel & Imray Northumberland House 303-306
High Holborn
London, WC1V 7LH(GB)

Description**BACKGROUND OF THE INVENTION:**

5 This invention relates to loop communications systems and more particularly to an adapter for coupling a host processor to a data communications loop, in which the adapter contains a local processor.

The token-ring local area network is a data communications system functioning to transfer data among a number of stations coupled to a one-way signal path. Each station has a host processor with memory and various peripheral devices to function as an independent workstation. Access to the ring is controlled by 10 token passing wherein a free or busy token is passed from one station to the next around the ring. An early example of the token ring is shown in U.S. Patent 3,597,549 issued to Farmer & Newhall and assigned to Bell Telephone Laboratories. In further refinement, the so-called Munich ring has evolved from developments by IBM, reported at Proceedings of the National Telecommunications Conference, Nov. 1981, pp. A.2.2.1-A.2.2.6. Standard protocols for these communications methods have been published by the IEEE as 15 the 802.5 token ring proposal, and a similar report by the European Computer Manufacturers organization.

It has been proposed to provide an adapter for connection between the host processor and the signal path of the ring. The adapter must operate relatively independently of the CPU of the host processor so that a minimum of interference with the functions of the host CPU is needed. The adapter must accommodate a variety of different types of host CPU, so that a special adapter need not be manufactured for each different 20 type of construction of the workstations. It is expected that entire office buildings or manufacturing facilities will be permanently wired to implement loop communications of this type between all computer terminals, word processors, telephones, bulk data storage stations, and the like. Thus, the adapters must be of relatively low cost, must be reliable and require little maintenance, must be of small size and dissipate little power, and above all must interface with a variety of differing systems.

25 In WO 84/03192, American Telephone and Telegraph Company, published on the 16th August 1984 and therefore comprised in the state of the art pursuant to Article 54(3) and (4) EPC, there is described a data network interface for controlling the transmission and reception of data packets to and from a signal path to which a plurality of stations are connected. The interface is provided to connect a host processor, forming one of the stations to the signal path and includes a local processor and a local memory with local bus means interconnecting them. The interface includes first and second memory access means respectively connecting the local memory and the host processor to the local bus means.

It is a principal object of this invention to provide an improved data communications system, particularly a communications adapter for coupling a host processor system to a token passing communications loop. Another object is to provide a communications adapter which has processing and control facility independent of the host system, and which operates with a variety of different host systems. A further object is to provide a low cost, low power, reliable and versatile communications adapter for token ring local area network systems or the like.

According to the present invention there is provided A digital data communications system comprising: a signal path, a plurality of stations connected to said signal path, each station receiving digital data from 40 the path and transmitting digital data to said path, each said station comprising a main processor containing a main CPU and a main memory, and having system bus means interconnecting said CPU and said main memory, a communications processor containing a local CPU and local memory means, and having local bus means interconnecting said local CPU and said local memory means, first memory access means connecting said system bus means to said local bus means whereby said main CPU can transfer data to 45 said local memory means and said local CPU can transfer data to said main memory, and second memory access means connected to said local bus means to access said local memory means for data transfer operations; characterized in that said communications processor further includes a timer having count register means connected to said local bus means and having a means to interrupt said local CPU when said count register means reaches a predetermined value, and bus arbiter means controlling access to said 50 local bus means, said first and second memory access means being operably coupled to said bus arbiter means, and said bus arbiter means receiving bus request signals individually from said first memory access means, said second memory access means, and said local CPU and granting access to said local bus means in a predetermined order of priority.

55 **BRIEF DESCRIPTION OF THE DRAWINGS**

The novel features believed characteristic of the invention are set forth in the appended claims. The

invention itself, however, as well as other features and advantages thereof, will be best understood by reference to the detailed description which follows, read in conjunction with the accompanying drawings, wherein:

5 Figure 1 is an electrical diagram in block form of a one way communications loop of the token-passing type which may employ a microprocessor device having features of the invention;

Figures 1a and 1b are timing diagrams of data frame formats used in the communications loop of Figure 1;

10 Figure 2 is an electrical block diagram of one of the stations in the communications loop of Figure 1, including a microprocessor CPU having features of the invention;

Figure 2a is a diagram of the word, byte and bit definitions of the 16-bit data word used in the local CPU in the microprocessor device 20 of Figure 2;

Figure 2b is a memory map for the local memory including local RAM 22 and ROM memory 27 of the local CPU device of Figure 2, and Figure 2c is a diagram of a number of adjacent memory locations in the general memory area referred to as a "workspace";

15 Figure 2c is a timing diagram showing voltage vs. time for operations in the local CPU of the microprocessor device 20 of Figure 2;

Figure 3 is an electrical diagram in block form of a MOS/LSI microprocessor chip including the local CPU or central processor unit employed as the microprocessor device 20 of Figure 2 and utilizing features of the invention;

20 Figure 4 is a greatly enlarged plan view of the semiconductor chip 20 containing the local CPU 21 of Figure 3 showing the physical layout of the various parts of the device;

Figure 4a is an enlarged detail view of a part of the layout of Figure 4 showing the register strip;

Figures 5a-5g are electrical schematic diagrams of various parts of the ALU and register strip within the microprocessor of Figures 3 and 4;

25 Figure 6 is a schematic diagram of a control ROM used in the local 21 CPU of Figures 2, 3 and 4;

Figure 7 is a state diagram showing conditions in the local CPU device of Figures 3 and 4 for an add operation;

Figure 7a is a diagram of a sixteen bit instruction word for the add operation of Figure 7;

30 Figure 8 is a schematic diagram of an example of a circuit within Figs. 3 or 5a-5g to be tested by microcode;

Figure 9 is a schematic diagram of the timer circuit.

DETAILED DESCRIPTION OF SPECIFIC EMBODIMENT

35

TOKEN RING LOCAL AREA NETWORK SYSTEM

Referring to Fig. 1, a communications loop of the token ring type is illustrated. A plurality of nodes or stations 10 are interconnected by a one-way signal path 11. A host processor system 12 is connected to each station 10. Each one of these systems 12 may be a computer workstation with a CPU, keyboard, disk memory, CRT display, and printer, for example. Signals on the loop propagate in one direction, as indicated by arrows, either bit-serial as will be described herein, or byte-serial, depending upon the requirements of the system. The purpose of this communications loop is to provide high-speed data transfer between processing systems. There is no central or master station in the loop; instead the loop operates on a peer-to-peer basis.

The communications loop of Fig. 1 is a local area network interconnecting stations 10 within a single office building, a building complex, or a campus, with the length of the signal path 11 being no more than a few kilometers. Using twisted-pair conductors for the signal path 11, data rates of about one to ten Mb/sec. are possible, and up to one or two hundred stations 10 may be connected in a single loop.

55 A control scheme based upon token passing is used for granting access to the ring of Fig. 1. A unique bit sequence, called a token, is passed from one station to the next. If a station has no data to transmit, the token is simply passed on to the next node. A station having data to transmit to another station 10 in the loop waits until the token is received, changes the token from "free" to "busy" and retransmits it, and transmits its data. After this transmitting station receives the message back at its input, confirming that its message has propagated around the loop, it retransmits the free token.

The communications loop of Fig. 1 operates according to the IEEE 802.5 standard for token ring local area networks. This type of system is described by Dixon, Strole and Markov in IBM Systems Journal, Vo.

22, Nos. 1 and 2, 1983, pp. 47-62; also, the system is described by N.C. Strole in IBM J. Res. Develop., Vol. 27, No. 5, Sept. 1983, pp. 481-498.

Each of the stations 10 has the capability of receiving data from the signal path 11 at its input 13, and transmitting data at its output 14. The data is transmitted on the twisted-pair signal path in the format of differential Manchester encoding. The data consists of a sequence of 8-bit groups or octets as illustrated in Figs. 1a and 1b. A free token consists of a one-octet starting delimiter, a physical control field which is two octets long, followed by an ending delimiter, as seen in Fig. 1a. The physical control field contains priority codes which allow levels of priority in granting access, as well as the "token", which is 1 for busy and 0 for free. The starting and ending delimiters contain bits which are code violations in Manchester encoding so that the delimiters will not occur in any address or data fields.

A data frame is shown in Fig. 1b, and consists of a starting delimiter, a two-octet physical control field containing a busy token in this case. The next field is the destination address, six bytes in length; a single loop would need only one byte for a maximum of 256 stations 10, but provision is made in the protocol for a much larger number of unique addresses. A source address field of six bytes follows, representing the address of the transmitting station. The information field is next, and this field is of variable length, depending upon the number of bytes of data to be transmitted; an average message is perhaps two or three hundred bytes of ASCII code, but thousands of bytes may be sent in one frame. The data is followed by a four-octet frame check sequence which contains a CRC-code for checking the addresses and data for errors; the receiving station performs this check. The frame concludes in an ending delimiter, followed by a one-octet physical control field which contains bits modified by the receiving station as the frame passes through to indicate that the address was recognized and the frame was copied, and also whether errors were detected in the transmitted data.

A transmitting station does not reinsert the free token onto the loop until it has received the beginning of its transmitted frame back at its input. As soon as it has recognized the starting delimiter, physical control field, destination address, and source address, it transmits the free token and strips the remainder of its transmitted frame. Each of the stations 10 other than the transmitting and receiving stations merely passes the frame on, but does not copy it. A station 10 can perform the following operations on the data stream:

- a) if the destination address is not that of the station, it retransmits the data frame without copying it;
- b) if the destination address is that of the station, it retransmits the data frame and also copies it;
- 30 c) the station can change the state of single bits in the received data (such as token or physical control fields) before retransmitting;
- d) a station can originate the transmission of data to another station;
- e) the originating station removes or strips a message frame from the ring that this station has previously transmitted after the message has gone all the way around the ring; this originating station retransmits a free token after it has passed on the starting delimiter, physical control field (including busy token), and addresses.

Referring to Fig. 2, each of the stations 10 includes a ring interface 15 which converts the incoming signal on line 13 to the voltage level for internal processing, and recovers the signal clock ϕ_5 from the loop signals. The ring interface also produces the outgoing signal on the output line 14 at the proper voltage/current level, based upon a transmit signal on line 14'. The incoming serial data on line 13' applied to a protocol handler device 16 which is a single-chip integrated circuit which implements the functions (a)-(f) listed above. The serial data is converted to parallel data in a register 17, and if the data is to be copied to parallel data in a register 17, and if the data is to be copied the bytes of incoming data are coupled by a FIFO buffer 18 to a data bus 19.

45 A message processor chip 20, another single-chip integrated circuit containing a local CPU 21 and a local RAM memory 22, receives the incoming data by DMA to RAM 22 via local address/data bus 23 and local control bus 24. The message processor also contains a general purpose timer 25 and a bus arbitrator 26, both coupled to the local bus 23 and control bus 24. The programming for the local CPU 21 of the message processor 20 is contained within a ROM 27 which is accessed by address bus 28 and data bus 50, extensions of the local bus 23; the ROM 27 is located on the protocol handler chip 16.

The incoming data, copied into the local RAM 22 by DMA, is then copied into the host processor 12 by DMA through a system interface chip 30, which is another single-chip integrated circuit.

The host processor 12 contains the host CPU 31, a main memory 32, and various peripheral devices 33 which would couple to the CRT, keyboard, disk drive, etc. A main address/data bus 34, and control bus 35, interconnect these elements. The host CPU 31 may be a part number 68000, or a 8086 device, as the interface device 30 accommodates the control and data formats of either.

The host processor 12 sends a message to another host processor 12 of the communications loop by first forming a message frame in its RAM 32 with destination address, its own address, and the bytes of

data to be sent as the message. This frame is copied into the local RAM 22 by DMA (or memory mapped I/O) through the system interface 30 and local bus 23. The message frame stays in the local RAM 22 until a free token is received and recognized by a decoder 36 in the protocol handler 16; when this occurs controllers in the protocol handler begin to fetch the frame from the local RAM 22 by DMA through local bus 23, data bus 29, and a FIFO 37, from which the data is serialized in a shift register 38 and goes out through a multiplex 39 to the output line 14'.

If the incoming data on the line 13' is not to be copied in this station, it does not reach the busses 29 or 23, nor the local RAM memory 22, but instead is applied through multiplexer 39 to the line 14'.

10

THE LOCAL CPU

Microprocessor devices with internal test of the invention to be described herein may be of various configurations; in this embodiment the microprocessor device takes the form of the chip 20 containing the local CPU 21 used in the system shown in Figure 2.

The microprocessor device 20 is a single-chip MOS/LSI device having the local CPU or central processing unit 21 as will be described in detail, along with local a read/write static random access memory 22. The CPU 21 and memory 22 communicate with one another by a 16-bit, parallel, bidirectional, multiplexed, address/data bus 23, along with the control bus 24. Suitable supply voltage and clock terminals are included; for example the device may employ a single +5V Vcc supply and ground or Vss, and a crystal may be connected to terminals of the device to control the frequency of an on-chip oscillator which will then define the local timing. It is understood that concepts of the invention may be used in various single-chip microcomputer or microprocessor devices with on-chip memory or off-chip memory, as well as in a microprocessor having separate address and data busses instead of the bidirectional bus 23.

In general terms, the local CPU 21 of Figure 2 functions in traditional manner. The CPU 21 fetches an instruction by sending out an address on the bus 23 and bus 28 to the ROM memory 27 then receiving the instruction via the bus 19 and bus 23 from the addressed location in the ROM memory. The CPU 21 then executes this instruction, which usually requires several machine cycles (as defined by the clock or crystal) including sending out via bus 23 addresses for the operands stored in the RAM memory 22, and receiving this data back on bus 23, as well as writing a result into RAM memory 22 by applying an address followed by data in successive cycles on the bus 23.

In the example to be described in detail, a 15-bit address is applied to the bus 23 from the CPU 21 (the 16th bit is hardwired 0 for addressing), which will directly address 64K bytes or 32K words of memory in the RAM 22 and ROM 27 as seen in the memory map of Fig. 2b. Each 16-bit word in memory includes two 8-bit bytes, so the CPU addresses 64K bytes.

The local CPU 21 of Figure 2 employs a memory-to-memory architecture in which the register files or working registers used by the CPU reside in the RAM memory 22. The RAM memory 22 is available for program or data storage, or for "workspaces" as seen in Figure 2c. The first word of a workspace is addressed by a workspace pointer (always an even address) and the remaining fifteen words are addressed by the workspace pointer plus 02 up to 1E (a 4-bit binary address).

Definitions of the functions of the control lines 24, the status bits for a sixteen bit status register ST, and interrupt levels are similar to Tables of U.S. Pat. 4,402,044, incorporated herein by reference. Of course, this specific microprocessor is merely illustrative of contemporary microprocessors or microcomputers in which features of the invention may be used.

The CPU 21 of Figure 2 executes the instruction set described in Table A and listed in Table B, although it is understood that other instructions could be defined for execution in addition to or in place of the particular ones described. Most of the instruction set of Tables A and B is the same as that of a microprocessor sold by Texas Instruments under the part number TMS9900 and described in a book entitled "9900 Family Systems Design", published in 1978 by Texas Instruments Incorporated, P.O. Box 1443, M/S 6404, Houston, Texas 77001, Library of Congress Catalog No. 78-058005; such book is incorporated herein by reference.

The processor architecture within the CPU 21 is shown in block diagram form in Figure 3, in chip layout form in Figures 4 and 4a, and in logic diagram form in Figure 5 and 6. Generally, the CPU includes an ALU or arithmetic logic unit 40 along with its associated registers, controls, data paths and related logic circuitry.

The ALU 40 consists of sixteen parallel binary adder/logic stages which function as the computational component of the processor. The ALU 40 receives a 16-bit input 10A and a 16-bit input 10B to produce a 16-bit parallel-output 10C. The input 10A of the ALU is from a 16-bit parallel A bus. The data on these internal busses is active low; the indicators A, A-, or NA should be used in referring to the bus data to be

technically accurate. The A bus may receive data from any one of several sources, with selection of inputs being made by microcode control inputs 41. Similarly, the input 10B is from a 16-bit parallel B bus which may receive data from any one of several sources as defined by control inputs 41. The output 10C from the ALU 10 goes to either a P bus or an E bus under control of microcode bits on lines 41. The ALU performs all the arithmetic and logic functions required to execute microinstructions in the CPU such as the functions of addition, subtraction, AND, OR, exclusive OR, complement, etc., as needed to execute the instructions of Table A.

The CPU 21 has a number of registers associated with the ALU 40, only three of which are accessible to the programmer. These three are a program counter or PC register, a workspace pointer or WP register, 10 and a status register ST. Other internal registers which are used during the acquisition or execution of instructions are inaccessible to the programmer.

The program counter PC is a 15-bit counter that contains the word address of the next instruction following the instruction currently executing. The fifteen stages are left-justified with the 16th bit hardwired to 0; instructions in the memory 27 are constrained to word boundaries, so a single byte is not accessed; 15 only even addresses (words) can be used. The CPU 21 uses the address in PC to fetch the next instruction from ROM memory, then program counter PC increments while the new instruction is executing. If the current instruction in the CPU 21 alters the contents of the program counter PC, then a program branch occurs to the location in ROM memory specified by the altered contents of the program counter. All context switching operations, such as interrupts, plus simple branch and jump instructions, change the contents of 20 the program counter. The program counter can be loaded from the E bus by lines EPC, or its contents applied to the B bus via lines PCB, or alternatively to the P bus via lines PCP. All such transfers to or from the program counter are defined by microcode control inputs 41, of course. Internally, the program counter PC includes binary add circuitry to add one to the count contained in the 15-bit register at the proper time, again under control of an input 41.

25 The status register ST is a 16-bit register that contains the results of comparisons, indicates status conditions, and supplies interrupt mask level to the interrupt priority circuits. Each of the sixteen bit positions in the status register signifies a particular function or condition that exists in the CPU 21; these bit position assignments are shown in the patent 4,402,044. Some instructions use the status register ST to check for a prerequisite condition; others affect the values of the bits in the register; and others load the 30 entire status register with a new set of parameters. Interrupts also modify the status register. All sixteen bits of the status register ST may be loaded from the E bus via lines EST or loaded onto the E bus via lines STE, upon a command on the control lines 41.

The CPU 21 uses blocks of words in its local RAM 22, called workspaces, for instruction operand manipulation instead of internal hardware registers. A workspace occupies sixteen contiguous memory 35 words as in Figure 2c, in any part of local RAM 22 that is not reserved for other use. The individual workspace registers may contain data or addresses, and function as operand registers, accumulators, address registers, or index registers. A large number of these 16-word workspaces may exist in the 32K words or the memory address space, providing a high degree of program flexibility. To locate the workspace in local RAM 22, the hardware register called the workspace pointer WP is used. The workspace 40 pointer WP is a 15-bit register (left justified with 16th or LSB hardwired to 0) that contains the memory address of the first word in the 16-word workspace as illustrated in Figure 2c.

The CPU 21 accesses any register in the workspace of Figure 2c by adding twice the register number R to the contents of the workspace pointer WP and initiating a memory fetch for that word. The 16th bit is zero so this is equivalent to adding two times the register number to WP. Figure 2c illustrates the 45 relationship between the workspace pointer and its corresponding workspace in memory. The WP register is loaded onto the B bus via lines WPB, or loaded from the DI bus via lines DIWP, under command of certain control lines 41 from the control ROM.

This workspace concept is useful during operations that require a context switch, which is a change from one program to another, as in the case of a subroutine or an interrupt. In a conventional multi-register 50 CPU, executing a context switch requires that at least part of the contents of the registers file be stored and reloaded, requiring a memory cycle to store or fetch each word. The CPU 21 accomplishes this operation merely by changing the contents of workspace pointer WP. A complete context switch requires only three store cycles and three fetch cycles, exchanging the program counter PC, status register ST, and workspace pointer WP. After the switch, the workspace pointer WP contains the starting address of a new 16-word workspace (Figure 2c) in local RAM 22 for use in the new routine. A corresponding time savings occurs 55 when the original context is restored. As seen in Table A, instructions in the CPU 21 that result in a context switch include: Branch and Load Workspace Pointer (BLWP), Return from Subroutine (RTWP), and an Extended Operation (XOP) instructions; these are in addition to device interrupts, the arithmetic overflow

interrupt, and others which also cause a context switch by forcing the CPU to trap to a service subroutine.

The internal registers not accessible to the user (programmer) in the CPU 21 include a memory address or MA register, a data or D register, and a K register. The D register is connected as input to a swap bus S via parallel lines DS, and a swap bus output on lines SK may be applied to the K register, all under control of microcode commands on the lines 41 from the control ROM. The D register functions principally as a data output latch, and is loaded from the E bus via lines ED. The D register is applied to a DI bus by sixteen parallel lines DDI. The data path from the output 10C of the ALU to the E bus and thus to the D register via lines ED, then to the DI bus via lines DDI and the A input of the ALU via the A bus is useful in divide routines, for example. Primarily, however, CPU output data is loaded into the D register from the E bus, then to swap bus S via lines DS, and then to sixteen address/data buffers 42 via lines 43, and thus to the local address/data bus 23. Data can be transferred onto the swap bus S straight or swapped, depending upon factors like byte operations being performed; these transfers are of course under control of microcode commands on lines 41.

The address/data buffers 42 are sixteen bidirectional, tristate buffers of conventional form, each having an input/output terminal connected to one of the local bus lines 23. These buffers usually receive addresses from the P bus via lines 44, or data via lines 43, for output to the bus 23; for input to the CPU 21, program or data words are applied via lines 43 to the swap bus S, thence to the K register via lines SK. It is also possible to load the P bus from the K register via lines PK, under a microcode command on a line 41, and thus output the K register via the P bus.

The addresses to local memory 22 or the ROM 27 are usually sent out from the CPU 21 via the P bus which is loaded by sixteen lines MAP from the MA register. The bits in this register can also be transferred to the B bus via parallel lines MAB, thus to the B input 10B of the ALU; alternatively the MA register may be loaded from the E bus via lines EMA or from the K latch via lines KMA, all as defined by microcode control lines 41.

Another internal register transparent to a user is the temporary or T register. This register receives a 16-bit parallel output 10F from the ALU 40, and applies its output to the B bus in three ways: either directly via line TB, shifted left via path TLB or shifted right via path TRB. The T register can also receive the B input 10B to the ALU delayed by 3/4 of a clock cycle by a path BT. The T register provides an important function in executing microcode for multiply and divide operations.

A register also used in multiply and divide operations is referred to as the MQ shift register (for multiplier/quotient). This register has the capability of right shift or left shift by microcode commands on lines 41. The register may be loaded from the A bus or the DI bus by 16-bit parallel lines AMQ and DIMQ, or may be outputted to the E bus or the B bus via lines MQE or MQB.

An instruction register IR provides the usual function of containing the current instruction, the instruction being executed during a given microcode state time (machine cycle). The instruction register IR is loaded from the DI bus via lines DIR, or may be loaded into the E bus via lines IRE, under microcode control via lines 41. Various fields of the instruction going to IR also can go to A bus by 2-bit and 4-bit connections IR2 and IR4. During each cycle, however, the contents of the instructions register IR are transferred via sixteen parallel lines IR0-IR15 to entry point and microcontrol generator circuits as well as miscellaneous control circuitry.

The microcode control signals 41 are generated in a control ROM 45 which is in this case split into two halves, 45H for the high side and 45L for the low side of the ALU and registers. Since there are many controls 41 used on only part of the bits, high or low, rather than all sixteen, space is saved by splitting the control ROM 45 in this manner. Each half of the control ROM has an X array 45X and a Y-select array 45Y. Buffers 45B for each of the lines 41 receive the select outputs from the Y array 45Y and introduce clocks or other logic as may be needed to produce the controls in the lines 41. The number of microcontrol lines 41 is about 130, depending upon the instruction set (Table A) to be implemented, well within the addressing range (256) of an 8-bit address on lines 46 that go to both sides 45H and 45L. The first of these 8-bit control ROM addresses on lines 46 is generated by entry point logic, and subsequent ones by a microjump circuit for executing a given instruction. Microjump addresses, produced on eight lines 47 which receive outputs 41 from the control ROM, can generate a jump address for the next state. The microjump address on lines 47 is fed back to a logic array 48 that can also generate an entry point from inputs received from an execute entry point array 49E or a source/destination address entry point array 49A. A group detect circuit 50 receives the 16-bit instruction word from IR as well as status bits from ST and other controls and produced two inputs to the entry point arrays 49A and 49E, first a group identification and second a field. The group is determined by the left-most 1 of the opcode as seen in Table B, and the field is three or four bits starting after the left-most 1. The address to the control ROM 45 on the lines 46 may also be held in an 8-bit latch 51 so the same state is re-executed as in multiply or divide instructions; to this end a 4-bit state

counter SC is provided which counts microcode state-times up to sixteen, and an overflow output of the state counter SC can control release of the holding latch 51. Thus, standard operation of the processor (as distinguished from test mode as will be described) is controlled by instructions loaded into the IR register to generate an entry point via group detect 50 and logic arrays 48, 49A, 49E; the entry point is a starting address for the control ROM 45 entered on address lines 46. This address results in a specific set of microcode commands on the control lines 41; some lines 41 will be active and most not. The address may also produce a jump address on lines 47 to define the control ROM address for the next microcode state, or the next state may be another entry point, or may be the same state due to the holding latch 51. When the last state needed for the instruction is reached, the next instruction is loaded into register IR and another entry point derived.

10 Interrupt codes received by the CPU 21 are applied to interrupt control circuitry 43. Bits 12-15 from the status register ST are also applied to the circuitry 43 to provide the interrupt mask for comparison with a interrupt code from external.

The control bus lines 24 are connected to control generator circuitry 44 which responds to lines 41 from the control ROM as well as to other conditions within the chip, and to the lines 24, to produce the necessary controls internal to the CPU and external to the CPU (on control bus 24).

A feature which aids in the test method is that the ALU 40 and its associated registers K, D, MA, PC, WP, T, MQ, ST and IR as described above are laid out on an MOS/LSI chip in a regular strip pattern as seen in Figure 4a. Each of these registers as well as the ALU contains sixteen bits or stages which are laid out in a pattern like memory cells, the bits of the registers arranged in parallel lines, and the busses perpendicular to these lines. The A, B, DI, E and P busses of Figure 3 are each sixteen parallel metal strips on top of the cells of the ALU 40 and registers (rather than beside the registers as depicted in Figure 3), and all of the dozens of control lines 41 are perpendicular to the metal busses and composed of polysilicon lines since they function as the control gates for transistors in the ALU 40 and registers of the CPU. The space needed in the control ROM 45 to generate the controls is about the same as the space needed for the controlled circuitry. This layout arrangement is advantageous when used with the "wide-word" control ROM or decoder of U.S. Pat. 4,402,043 because it fits alongside the control ROM with little wasted space for routing conductors. That is, almost all bus lines A, B, DI, E and P etc., and all control lines 41 are routed over or under functional regions or cells of the chip rather than over unused silicon, and almost all 90 degree turns are produced inherently at functional cells rather than in conductor routing. The enlarged view of Figure 4a shows a small part of the ALU 40 and registers, illustrating the regular pattern of metal bus lines and the polysilicon control lines 41 for an N-channel silicon gate MOS device made generally according to U.S. patent 4,055,444, assigned to Texas Instruments.

It is significant to note that most of the connecting lines such as KDI, ED, EMA, MAB, DINP, etc. as mentioned above are not physically lines or elongated conductors at all but instead are merely metal-to-silicon contact areas along the metal bus lines of Figure 4a. That is, routing of 16-bit parallel conductors is minimized by the strip feature.

40 THE ALU AND REGISTER CIRCUITS

The detailed circuits employed in the ALU 40 and its associated registers will be described with reference to Figures 5a-5g. These circuits fit together as seen in Figure 3. Generally, only one bit of the sixteen bits is shown in a strip. For the most part the remaining fifteen bits are the same as the one shown in detail, with some exceptions.

The ALU

50 Referring to Figure 5a, one of the bits of the ALU 40 consists of a complex logic circuit having its input 10A connected to the A bus and its inputs 10B connected to the B bus while its output 10C is connected by transistors 10a and 10b to the E and P busses, respectively. The transistors 10a and 10b have connected to their gates control lines 41 labelled H1ALIE and H1ALIP, respectively. In this description, the convention used for labelling commands or control lines 41 is (a) clock phase such as H1 or "half 1", then (b) the source such as "AL" (for ALU), and (c) "t" or "I" for to or from, followed by (d) the destination such as "P" or "E" bus. The half 1 and quarter clocks are shown in the timing diagram of Figure 2d. Using this convention, H1ALIP means this control 41 occurs during the H1 or half-1 clock and defines the connection from ALU to the P bus. An input to the A bus from the DI bus by line DIA and transistor 10c is controlled by

an H1DltA command on a line 41, translating as "DI to A, on H1 clock". Carry-in and carry-out for this bit are on lines Cout and Cin respectively, which are connected to adjacent stages of the ALU. The carry-out line is precharged to 1 or Vcc on Q1 by transistor 10d, then conditionally discharged after Q1 goes low by a transistor 10e. A NOR gate 10f drives the gate of the transistor 10e based upon a S1STOPG input on one of the control lines 41 (meaning stop generate, on S1), and the voltage on a node 10g. The output of gate 10f is the carry generate condition, and is made unconditionally 0 if STOPG is active. The node 10g, precharged to 1 on Q1, is connected to a node 10h by a transistor 10i which has the inverted A input on its gate. Similarly, in the carry propagate circuit a node 10j is precharged to 1 on Q1 and is connected by a transistor 10k (also receiving inverted A input) to a node 10m. The nodes 10h and 10n, precharged to 1 on Q1, are conditionally discharged by complex NAND/NOR logic circuit including transistors 10o driven, by the B and B input (inverted and twice-inverted), along with transistors 10p driven by the four control lines 11 labelled ALU1 to ALU4 (all valid on H2). The four controls ALU1-ALU4 define the operation performed in the ALU as set forth in Table C. The "H3Logic" control on a line 41 is the arithmetic/logic control; when this line is high the ALU performs logic functions by unconditionally grounding the Cin line by transistors 10q, but when low Cin is applied to an input of an exclusive NOR circuit 10p, the other input being the inverted propagate signal from node 10j. The propagate signal also drives the gate of a transistor 10s in conventional manner.

A carry input to the LSB of the ALU 40 is produced by microcode controls 41. A ClCO or "carry-in from carry-out" control applies the carry-out from the MSB of the prior state to the carry-in at the LSB for the current state. A STCIN control applies the status bit-3 to the carry-in. A CIN control produces an unconditional carry in to the LSB.

The Workspace Pointer Register

25 A detailed circuit diagram of the workspace pointer register WP is shown in the lower part of Figure 5c. This register includes two static inverter stages WP_a with an input node WP_b connected to the output for feedback during H2. The input node can be loaded from the DI bus by a transistor WP_c having an H4WP1DI (WP from DI, on H4) command on its gate. An intermediate node WP_d is connected to the B bus via an inverter WP_e and a transistor WP_f which has a command H1WP7B (WP to 8 on H1) on its gate. Fifteen bits are exactly the same, and bit sixteen is a hardwired 0 or V_{ss}.

In addition to functioning as the workspace pointer, the register WP is used as the B input to the ALU 40 when generating the source and destination addresses for typical instructions of Table A. In a context switch, WP is saved, by writing into memory 22, and the path for this operation is the B input of the ALU 40 35 and the P bus. WP is loaded from off-chip by the DI bus.

The Program Counter and Incrementer

40 Also seen in Figure 5c, the program counter PC consists of a pair of static inverters PC_a having an output node PC_b connected to an input node PC_c on Q2 by a transistor PC_d. The input node PC_c may be loaded, from the vertical E bus by a transistor PC_e having H4PCIE (PC from E, on H4) on its gate; this signal is on one of the control lines 41 running horizontally through the register matrix. The output of the program counter at an intermediate node PC_f is connected to the gate of a transistor PC_g, from which the program counter contents may be read out onto the vertical P or B busses by transistors PC_h or PC_i and lines PCP or PCB. The signals to control these transfers are H1PCIP (PC to P, on H1) and H1PCIB (PC to B, on H1) on separate horizontal control lines 41. The program counter is incremented when an H3PCINC signal appears on one of the control lines 41, turning on a transistor PC_j; on H1 a transistor PC_k turns on, loading the contents of this bit of the program counter onto the gate of a transistor PC_m in series with a carry line PC_n from the prior bit. Each bit of the carry path is precharged to 1 on Q1 by a transistors PC_p. The carry-in for each bit on line PC_n is connected by a NOR gate to the gate of a transistor PC_q; the NOR gate also has the H4PCIE signal as one input. This circuit causes a carry to be propagated if the bit is 1 and carry-in is 1; or causes the bit to go from 0 to 1 if the carry-in is 1. The LSB stage is a hardwired 0 because only "the 15-bit word address is sent out on the bus 4; the 16th bit is the byte address which is not used for access. For byte operation using the odd numbered byte, the next lower even number is the address then the byte swap circuitry S is activated. Thus, the 16th bit of the address is always 0. All the other fifteen bits are identical to that shown in Figure 5c. The carry-out from the last bit is truncated as it would represent address FFFF (in hex).

The Memory Address Register:

One bit of the MA register is seen in detail in the upper part of Figure 5c. This register includes a pair of standard inverters MAa and MAb in each of the sixteen bits with feedback by a transistor MAC clocked on H2. The register is loaded from the E bus via sixteen separate lines EMA and a transistor MAD in each line. The gates of the sixteen transistors MAD are driven by a signal H4MAIE (MA from E, on H4) on one of the horizontal control lines 41. Likewise, each MA register bit is loaded from the K latch by a lines KMA and a transistor MAe which has H4MAEK on its gate. Output from the MA register is taken at an intermediate node between the two inverters which is connected by a transistor MAF to two output transistors MAG and MAh. An H1MAIP command on one of the control lines 11 turns on the transistor MAG and connects the output to the P bus via line MAP. An H1MAIB command turns on the transistor MAh and connects the output to the B bus via the line MAB. Of course, it is understood that there are sixteen of each of the input and output lines to or from the busses, as is true for the other registers.

The MA register is usually loaded via P bus to the A/D buffers 12 as a basic part of most instruction sequences. It is usually loaded from the ALU output 10C via the E bus when an address is generated by adding WP + 2S, for example, which is usually done for most instructions.

The Data Register:

The D register or data register contains sixteen bits constructed as shown in detail in Figure 5b. This register stage consists of a pair off inverters Da and Db with a feedback path via transistor Dc clocked on H2. The input of this register is from the E bus via transistor Dd and a line E (one of sixteen lines ED), with the transistor Dd being controlled by an H4DIE (D from E) command on one of the lines 41. One of the 25 outputs is from an intermediate node De and a transistor Df which connects to the DI bus by a transistor Dg and one of the sixteen lines DDI. This output is controlled by an H1DtDI (D to DI) command on a control line 41 which is connected to the gates of all sixteen of the transistors Dg. The other output from the D register is by sixteen lines DS, each going to he gate of a transistor Sa in the swap circuit S. The output of this transistor inverter is connected to the swap bus Sb by a transistor Sc having a "straight" command on its gate or a transistor Sd having a "swap" command on its gate. The line 41 connected to the gates of each of the sixteen transistors Sc carries the H2DtSS (D to swap-straight) command, while the line 41 carrying the H2DtSW (D to S, swapped) is connected to the gates of the transistors Sd. E bus is precharged to Vcc via Dh and Dd.

The D register most often functions to receive the data output from the ALU 40 via the E bus. Also, it functions as the source of an ALU operand input to the A side via the DI bus and the A bus. The D register is used mainly for data output, and also in ready and hold conditions and in divide operations, for example.

The K Latch:

The K register or K latch contains sixteen identical stages, one of which is shown in Figure 5b. The K register uses two standard inverters Ka and Kb with feedback on H1 by a transistor Kc. The output is connected to the gate of an inverter transistor Ke which is connected to output transistors Kf and Kg. A command Q1KIP (K to P) on one of the lines 41 turns on the transistor Kf and connects the output of the K register to the P bus via line KP. A command Q1KIDI (K to DI) turns on each of the sixteen transistors Kg and connects the output to the DI bus by sixteen lines KDI. The K register is loaded from the swap bus S in either straight or swapped condition by transistors Ki and Kj. The command H3KfSS on one of the control lines 11 connects the swap bus to the input of the K register via sixteen transistors Ki while a command H3KfSW connects the swap bus to the input of K via sixteen transistors Kj for swap or byte operations.

50

The Swap Circuit:

The swap circuit S shown in Figure 5a functions to connect the 16-bit D register to, or the K register from, the A/D buffers 42 via lines 43, either straight or with the high and low bytes swapped. Addresses are usually transferred in or out via the P bus that is connected to the A/D buffers without going through the swap bus, since addresses need not be swapped in byte operations, only data. Thus, data comes in via the K register and goes out via the D register, both using the swap bus S.

The Temporary Register:

In Figure 5d one stage of the sixteen stages of the T register or temporary register is seen in detail. This register consists of two standard inverter stages Ta and Tb with feedback on H2 via transistor Tc. Output from the T register is from a node Tf at the output of the inverters, via an inverting transistor Te to a node Tf precharged to Vcc on Q3 via transistor Tg. A H1NLDI command on one of the lines 11 controls a transistor Th in series with transistor Te. Output from the node Tf directly to the B bus is by a line TB and a transistor Ti controlled by a H1TIB (T to B) command on a horizontal line 41. To shift left, the node Tf is connected by a transistor Tj to the B bus in the next more significant bit, to the left, via line TLB. To shift right, the node Tf is connected by a transistor Tk to the B bus in the next lower significant bit, to the right, via line TRB. The transistors Tj and Tk are controlled by commands H1TSLB and H1TSRB on two of the horizontal lines 41. The T register is loaded from the B bus by a line BT and a transistor Tm clocked on Q3 along with an inverter Tn and a transistor Tp which has H1TIB on its gate. This input is delayed by 3/4 of a clock cycle. The transistor Tp at the inverter output has its output to the input node Tq to load the T register. The transistor Tr at the inverter output and a transistors Ts with H1TIB on its gate goes to node Tf as a path for replacing the B bus information on the B bus after a delay. Thus, H1TIB followed by H1TIB is a quicker path than loading T register then reading it out onto the B bus in the next cycle. The "F" output from the ALU 10 is connected to the input Tq through a transistor Tt, and to the gate of a transistor Tu. The transistor Tu is in series with a transistor Tv, and both Tu and Tv are controlled by H1TTF. The output of the transistor Tv goes to the node Tf as an output from F to the B bus, which may be either straight, left-shifted, or right-shifted.

The MQ Shift Register:

The MQ register consists of sixteen bits, one of which is shown in Figure 5f. This register may be shifted left or right by controls 41 and so is used in multiply and divide instructions. Also, the register can be used as a general purpose working register. To this end, an input node MQa may be loaded from the E bus via transistor MQb and line EMQ by control H4MQIE, or loaded from the DI bus via transistor MQc and line DIMQ by control H1MQIDI. Feedback on Qe is provided by a transistor MQd, and the output of the first stage is connected to the input of the second during H2 by a transistor MQe. The left shift function is provided by a transistor MQf connecting the input node MQa to a node MQg in the next lesser significant bit of the register, this occurring when a command HD4MQSL appears on one of the lines 41. The HD4 prefix for this control means that it occurs at H4 in the next state time, or delayed one clock cycle, from the time this microcode is generated in the control ROM 45. The right shift function is produced when an HD4MQSR control occurs on a line 41 turning on a transistor MQh to connect the node MQa to the input node MQa of the next more significant bit of the register. The contents of the MQ register are applied to the E, P or B busses by a circuit consisting of an inverter MQi and inverter transistor MQj connecting input node MQa (twice inverted) to output node MQk. The output node MQk is connected by transistors MQm to the B, P and E busses when H1MQIB, E1MQIP or HD1MQIE commands occur. The output MQn of the first inverter MQp is connected to the input of the second stage by a transistor MQq on H4MQIE, when MQb is activated.

The Status Register:

One bit of the sixteen bit status register ST is shown in Figure 5g. This stage consists of a pair of inverters STa and STb, with feedback on Q4 by a transistor STc. An input node STd may be loaded from the E bus through a transistor STe when a command H1STIE occurs. Several other input transistors STf may set or load the input node from other sources, such as other controls 41 from the control ROM 45, etc. The output of the status register is taken at an intermediate node STg by an inverting transistor STh with a transfer transistor STi going to the E bus, under control of an HD1STIE command on a line 41.

Among the controls 41 from the CRROM 45 are twelve control status signals CS1 to C12 and two set status signals SS0 and SS2. These control various transistors STi to set or conditionally set the status bits according to the instruction set of Table A and the status bit definitions.

The Instruction Register:

The instruction register IR is a sixteen bit register, one bit of which is shown in Figure 5g, containing two inverter stages IRa and IRb with feed back via transistor IRc on H2. The instruction register is loaded from the DI bus at input node IRd through transistor IRc on an H4IRLD command. Also, the instruction register may be cleared by a transistor IRf connecting the input node IRd to Vcc on an H1IRCLR command.

5 True and complement outputs from the instruction register IR are taken at nodes IRg and IRh; these outputs go to the group detect 50 as the IRO-IR15 (and IRO-IR15) signals, and are of course used to generate entry point addresses, etc. Coming in, bits which define the addresses of registers in the workspace are fed directly to the A inputs of the ALU for source and destination address generation, under control of command S 41 before reaching IR.

10 Certain ones of the IR outputs and ST outputs are used in a jump detect circuit 56 (Figure 3) to detect any of the jump instructions and/or conditions as defined by the instruction set of Table A. The bits received by the jump detect circuit are: ST0-ST3, ST5, ST0-ST4, IR4-IR7, and IR4-IR7.

The group detect 50 receives all fifteen IR bits and complements and determines which of the eleven groups the instruction word falls in. This is based on the position of the leading "1". Then, the following four 15 bit (or in some cases three bit or two bit) field of the instruction is used to generate the entry point address in a PLA.

The Compressed Control ROM:

20 The microcode control signals on the lines 41 which define the operation or the CPU are generated in a control ROM 45 as seen in Figure 6. According to U.S. Pat. 4,402,043, the control ROM is compressed so that it uses much less space on the chip, and is wide-word format so control bits are easily added, such as those used for the test modes. The control ROM 45 generates a different set of microcontrol signals on the 25 lines 41 for each individual address applied to the eight input address lines 46; the address input is split into a four-bit X address on lines 46a and a four-bit Y address on the lines 46b. For each of the 256 possible address inputs, a unique combination of outputs could be produced, but in a typical embodiment less than 256 are required because an instruction set can be implemented with fewer than 256 states. In a CPU which executes the instruction set of Table A, for example, about 150 addresses or microcode states 30 are used to selectively activate about 120 to 130 controls 41 (including eight microjump addresses for lines 47).

The control ROM 45 is split into an X-select portion 45x and a Y-select portion 45y. The X-select portion contains sixteen X lines 45a and a variable number of Y lines, depending upon the degree of compression. The X address on the lines 46a selects one-of-sixteen of the lines using a standard decoder.

35

MICROCODE STATES FOR REGISTER-TO-REGISTER ADD OPERATION:

40 One of the most common instructions in any processor is "add". In the CPU described herein, the add instruction uses workspace addressing. The two operands are obtained from registers in the local RAM 22 via bus 23 and the result is stored in a register in the local RAM 22. Various addressing modes may be used, including direct, indirect, automatic incrementing, indexed, immediate or relative, as explained in Pat. 4,402,044. The example uses register direct addressing for one operand located in the workspace, and register indirect addressing for the other operand which is in a register whose address is located in the 45 workspace; the result is stored in the same register which is directly accessed for an operand. This instruction is referred to as "ADD "R, R" in assembly language for programming purposes.

45 Execution of the add instruction of the example uses parts of seven "state" times or clock cycles. A different set of microcode controls is produced by the control ROM during each microcode state time. The timing sequence of clocks, used in the CPU 21 is shown in Figure 2d and includes four overlapping half-cycle clocks H1, H2, H3 and H4. The clock H4 of course occupies part of the state time after the one in question. Also, a state time contains four quarter cycle clocks Q1, Q2, Q3 and Q4. The device is built using N-channel technology, so the clocks are positive-going. Typically, the state time is 200 nsec., or the repetition rate is 5 MHZ. The clock input 0 to the chip or the crystal frequency is four times the state frequency or about 20 MHz.

50 The local RAM 22 can be accessed within a state time of the CPU 21, so when an address is sent out on the bus 23 during Q2 of a given state time the contents of the addressed location will appear on the bus 23 for input to the CPU 21 during Q4 or H4 in the same state time. The add instruction of the example requires five memory access operations: one to fetch the instruction, one to fetch the address of the

"source" operand, two for the "source" and "destination" operands, and one to store the result in the local RAM 22 at the "destination" location.

The instruction for the add operation of this example is fetched before the previous instruction has completed execution. This pre-fetch increases operating speed. It will be assumed that the prior instruction also used an ALU cycle and a destination write, like an add operation, and that the next instruction following the one of the example is also similar to the add operation, to illustrate the pre-fetch and overlap of operations. Of the seven microcode state times partially occupied by the add operation of this example, the first two are shared with the prior instruction and the last two are shared with the following instruction execution; thus only the central three are exclusively used by this add operation.

10 Referring now to Figure 7 a chart of microcode states for execution of the ADD "R, R" instruction is illustrated. The seven state times to be discussed are designated S1 through S7. Briefly, during S1 (which overlaps the prior add instruction) the instruction word in question is pre-fetched by sending out the address which is in the P register onto the bus 23 to access a location in the memory space (ordinarily the ROM 27 which contains program rather than the RAM 22 which is used for workspaces and data); then the contents 15 of this location are received back on the bus 23 and loaded into the instruction register IR. During S2, the result of the prior operation is written into the location in the memory 22 which is in the memory address register MA; the contents of MA register are sent out on the bus 23 followed by the result of the previous add operation. Meanwhile, during S2, the contents of the WP register are added in the ALU 40 to a number derived from the instruction word pre-fetched during S1; this number is two times the source register 20 number, or 28. So, the address of the source register is calculated during S2 and is stored in the MA register, then during S3 this address is sent out on the bus 23 to access the source address contained in the selected register of the workspace; this source address is read from the local RAM 22 and returns via bus 23 to be stored in MA and K registers in CPU 21. During S4 state time the address of the source operand is sent out from the K register onto bus 23 and the contents at this address in memory 22 returned 25 via bus 23 to the K register.

Also during S4 the destination address is calculated in ALU 40 by adding the contents of the WP register to two times the "D" field of the instruction word and storing the result in the MA register. Next, during S5, this calculated D address (within the workspace, thus "direct") is sent out from MA register on bus 23 while the operand in the K register is moved to the T register. Then, the contents of the addressed 30 register in memory 22 is returned via bus 23 to be stored in the K register of the CPU 21, still during S5. Now the add operation is ready to be implemented; the S operand is in the T register and the D operand is in the K register. Thus, during state S6 the add operation is executed by adding the contents of K and T and putting the result in the D register. Meanwhile, however, during S6 the instruction for the next operation is pre-fetched by sending out the address in the program counter PC (which has been incremented) onto local bus 23 and receiving the contents of this address from ROM 27 back via bus 23 to be loaded during 35 S7 in the instruction register IR in CPU 21. Also, during S7 state time, the result of this example add operation is written into memory 22 by first sending out the address in the MA register (which is the D or destination address calculated during S4) onto bus 23, followed by the contents of the D register, completing execution of the ADD "R, R" instruction. During S7, incidentally, the source address for the next 40 instruction is calculated in the ALU 40 by adding the contents of the workspace pointer WP to two times the S field of the instruction word accessed in S6, so another add operation is already two state times into its execution sequence. Accordingly, the states S6 and S7 correspond to the states S1 and S2 of Figure 7.

The instruction word for "ADD R, R" is shown in Figure 7a. This is the sixteen bit word which is read during S1 state time in Figure 7, i.e., the contents of the location addressed by the PC register, in this 45 example. The first 3-bit field 010 says "add". The "B" field, bit 3, defines whether this is a byte operation or a word operation; in this example it is a word operation. If B = 1, the operands are bytes and the operand addresses are byte addresses. With B = 0, the operands are words. The TD and TS fields, bits 4, 5 and 10, 11, determine the addressing mode of that operand. In the example, TS is 00 so the S field (bits 12-15) contains the register number in the workspace which has the source operand. TD is 01, meaning indirect 50 workspace register addressing mode, so the D field (bits 6-9) contains the register number in the workspace which has the address of the destination operand in it.

Considering the execution of the ADD "R, R" instruction example in more detail, it will be noted that Figure 7 also shows the contents of the various registers and busses as a function of time, as well as other features. Each of the register and bus operations will be examined for each microcode state time, along 55 with the control signals produced on the lines 41 to produce these operations.

IN the S1 state time of Figure 7, the control line 41 for H1PCIP (see Figure 5c) is high during H1, turning on transistors PCI and placing the contents of the PC register on P bus. Then the DEN command goes high on H1, so the P bus is loaded into the output buffers 42 and thus to the bus 23. The program

counter PC is incremented during H3 time by the H3PCINC command on a line 41, turning on transistors PCj and PCm of Figure 5c, so later at S6 the next instruction in sequence will be accessed. A DEN signal comes up on H2 on a control line 41 to produce a DEN-command at H3 on one of the lines 24 to enable the memory 22 to put data on the bus 23 beginning at Q4. The instruction word fetched here is valid on the bus 23 beginning at Q4, and is loaded into the K register via lines SK by a control H3 KISS generated every H3 except when blanked; this control turns on the sixteen transistors Kj of Figure 5b. During this state time S1 an add operation for the prior instruction occurs just as will be described for S6, directing the result to the E bus. A control SAMPI on one of the lines 41 causes the interrupt input to be sampled so that if an interrupt is present a context switch will occur.

10 Turning now to the S2 state time of Figure 7, the instruction which is in the K register is connected to the DI bus at Q1 when the Q1KjDI command goes high and turns on the sixteen transistors Kg; this command occurs every Q1 unless blanked by a control line 41. A command 2STA at H1 causes the S field of the instruction word on the DI bus to be left-shifted and connected by four transistors Ca to bit-11 through bit-14 of the A bus as seen in Figure 5e. The contents of the workspace pointer register WP are transferred to the B bus at H1 by the H1WPB command on a line 41 turning on the transistors WPf of Figure 5c. Thus, with WP on the B bus and 2S on the A bus, when the ALU 40 operates (at H2 and H3) an output will be produced at node 10C during H3 which is the sum (WP + 2S). A command ALTE on a line 41 occurs at delayed H1 time, meaning delayed one state time, so at H1 of the next cycle the ALU output is connected by transistors 10a to the E bus. As will be described below with reference to the S7 state, the result of the previous operation is written into memory 22 from D register which was loaded from the E bus at H4; the ALU 40 loaded the E bus at S2, H1 (the same as H1 for S1).

During the S3 state time of Figure 7, the address of the register which will contain the source address is generated by an ALU operation. In this operation the contents of the workspace pointer register WP are applied to the B bus by an H1WPB command turning on transistors WPf and an IR2D command at H1 which applies bit-6 through bit-9 of the instruction register IR to bit-11 through bit-14 of the DI bus via the transistors IRf of Figure 5g. The IR2D command in effect left shifts the D field of the instruction word (see Figure 7a) by one bit to multiply by binary two, then applies it to the A input of the adder via the DI bus. The ALU 40 is in the add condition by default, none of the ALU1-ALU4 commands being present, and the input 10c is applied to the E bus at H1 of the next state time by the H01ALTE command on the line 41 to the gates of transistors 10a (Figure 5d).

TESTING OF CPU

35 As thus far described, the normal operation of the CPU 21 in executing the instruction set of Table A has been illustrated. For testing the chip containing the CPU 21, however, it is necessary to add some further capabilities. It is too time consuming, and not totally conclusive, to test a one-chip CPU device by executing all of the instructions, because some faults are data-dependent. Also, stuck-faults, i.e. internal nodes stuck at logic-1 or logic-0, may cause an erroneous result only under conditions occurring rarely. So, 40 as in testing a dynamic memory device, all of the data bits in the registers of the CPU, and nodes in the ALU, must be tested for various patterns of data and controls.

This testing is accomplished in the CPU 21 by employing additional microcode in the control ROM 45 which has no function other than test, and which is not used in the execution of the standard instruction set of Table A. This test microcode may use unique control lines 41, and microstates, and also will use sets of signals 41 like many existing microcode states for the standard instruction set, but different microjump addresses will be needed because the sequences will always be different from the standard sequences like Figure 7.

Each of the register bits of the registers in the CPU may be thought of as a memory element ME of Fig. 8, with an input MEI and an output MEo. The input MEI can receive a data bit from two sources, IN-A or IN-B, and these input paths are controlled by the control nodes IN-1 and IN-2 which connect to gates of transistors. The output MEo may also be connected to more than one output bus, as seen for the MQ register which can go to P bus or B bus, or the K latch to the P bus or DI bus. However, for simplicity in explanation of the fault-flushing concept, only the inputs will be treated. For example, the general form of Fig. 8 corresponds to the two inputs MAIDI and MAIE via transistors MAe and MAd to one bit of the MA register of Fig. 5c, where bits of the DI bus and E bus would correspond to the sources IN-A and IN-B.

A truth table for the circuit of Fig. 8 is shown in Table D. The patterns T to Z are various combinations of inputs data bits IN-A and IN-B with controls IN-1 and IN-2, showing the memory output MEo for the initialization cycle I(T) and for the test cycle M(T+1) for a good circuit and for a faulted circuit. Each of the

fault possibilities are shown for each input; these are stuck-at-one S@1 and stuck-at-zero S@0. Note that all of these patterns are needed to find all of the possible faults. Pattern-T finds stuck-at-one for either control, IN-1 S@1 or IN-2 S@1. Pattern-U and pattern-W find stuck-at-zero for the control bits, IN-1 and IN-2. Pattern-X and pattern-Z find stuck-at-one for the data bits IN-A and IN-B. Pattern-V and Pattern-Y finds stuck-at-zero for the data bits, IN-A and IN-B.

To implement the test of Table D in the registers of the CPU 21, algorithms are written in microcode, for each register. For example, in the first state the memory element (actually all sixteen bits of a register at once) is initialized by providing a condition of all input data high; this is done by inputting FFFF via the bus 23 with DEN active, and selecting microcode which places this data on the two busses or sources for the selected register, e.g. K register and E bus for MA register, or DI bus and E bus for MQ register bits. The controls for the two inputs are held high at this point, e.g. by KIMA and EIMA microcode bits 41, or DlMQ and EMQ microcode bits 41. The memory contents are verified to be high by either (1) outputting this register via P bus and bus 23 for external check by the test machine, or by an ALU operation and jump or interrupt on status depending upon the results. This completes the initialization cycle (I1). Next the pattern-
10 T is produced; all controls are held off by the appropriate microcode bits 41, and the input data (the busses, etc.) are held low by microcode controls or data input of zeros. The memory content is checked at this point, as before, and the result interpreted as in Table D, detecting if either control is stuck-at-one. This sequence can be repeated for each of the registers, and indeed for the A,B, DI or P busses themselves.
15 The difference from ordinary operation is that two inputs are simultaneously trying to drive a node in the test microcode, whereas this would not be used in executing the standard instruction set, and indeed would be highly undesirable. However, this condition is needed for stuck fault detection. Errors would certainly turn up in ordinary operation if stuck faults existed, but these would be data-dependent and may require exhaustive and time-consuming patterns on the test machine.

For the other patterns U,V,W,X,Y,Z, tests are made by first initializing as described above, and verifying
20 that all memory bits are high. Then the selected pattern is produced by inputting data to the two selected inputs for the selected register according to the values given in the truth table, and also defining the two controls according to the truth table, and verifying that the memory is low or high as set forth. This is iterated until all possible combinations are checked. Generally all bits at an input are one or zero, at a given time. Doing this same sequence for all registers of the CPU increases the number of iterations in the loop at
25 this point, but does not add to the number of microcode states since merely the input data is changed. These patterns check for controls stuck-on-zero, or data bits stuck-on-zero, or data bits stuck-on-zero.

One additional test of the CPU 21 is a check of the outputs 41 from the control ROM 45. For this test a shift register SR is included on the chip at the control ROM output, and this shift register is loaded in parallel by the control ROM 45 output by a particular microcode bit #LDSR. In standard operation the outputs 41 are not loaded into SR, but for this test the #LDSR bit occurs and all 130 (or whatever number) bits are loaded into SR in one cycle, then for the next 130 cycles the bits are clocked out serially on one of the output pins, such as an address/data I/O pin of the bus 23. If there are 256 possible addresses for the control ROM 45, then this entire sequence must be repeated 256 times. Addresses for the control ROM are applied to lines 46 from a counter TC, which is an eight-bit shift counter, for example, zeroed at reset and incremented every time the microcode bit #LDSR is produced. Thus, a selected code input on the four test pins (TEST, TEST0, TEST1 and TEST2) as defined in Table E causes the test control circuits to produce the control #LDSR which increments the counter TC, applies the counter output to the address bus 46, producing a set of microcode outputs 41 from the control ROM 45, which is loaded into the shift register SR (rather than executed in the CPU); then during the next 130 or so clock cycles this microcode word is shifted out serially one bit at a time, to be compared with the known value in an external test machine.

ON-CHIP TIMER

50 Referring to Fig. 9, the timer 25 is shown in block diagram form. The timer includes a fixed 512-bit prescale counter 25a which receives the local clock at its input 25b and produces a decrement signal at its output 25c. The prescale counter 25a may be constructed as a pseudorandom shift counter with exclusive-OR feedback, as is well known. The output 25c is produced about every 170 microsec, and goes to a timer data register 25d which may be loaded with any 8-bit value by an input 25e from a timer latch 25f. The latch 25f may be loaded from the local bus 23, and it appears in the memory space accessed by the local bus. The latch 25f can also be read by the CPU 21 via the local bus 23 through path 25g. In like manner, the contents of the timer data register 25d can be read by the local CPU 21 via the local bus 23 through path 25h, as this register also appears in the memory space accessed by the local bus. Every time the

prescaler 25a decrements to zero and produces an output 25c, the register 25d is decremented, so when a count of the value loaded into this register 25d from the latch 25f and local bus has been reached, the register 25d will reach through an all-zero condition. This condition is detected by a zero-detector circuit 25i, which is a NOR gate arrangement of conventional form. An output 25j from the zero-detector circuit is connected to the set input of a set-reset flip-flop 25k, which receives its reset input 25m from a certain bit of the MPSTS control register. If the flip-flop 25k has been reset, then when it is set by input 25l it produces an output 25n which is used to indicate a timer interrupt to the CPU 21. An AND gate 25p receives this interrupt signal as one input, and receives a bit from the MPCTL register on input 25q which determines enable or disable of the timer interrupt. A "timer interrupt active" signal is thus produced on the line 25r going to the CPU 21, if all conditions are met.

The GPTIEN bit in the MPCTL register and on line 25q is set or reset by the CPU 21 by writing a word to MPCTL. When set, this GPTIEN (general purpose timer interrupt enable) bit enables a interrupt (priority level ten) when and if the GPTINT (general purpose timer interrupt) bit is also set in the MPSTS register on line 25m. If GPTIEN is zero, then GPTINT may be set, but no interrupt is generated. The GPTINT bit is set on the local bus cycle following the decrement pulse of register 25d from a count of one to a count of all zeros. This bit is automatically set by the circuitry, and is reset by the CPU 21 under its program control. It is set when the timer decrements past zero, then is cleared when the CPU 21 writes a word to MPSTS register containing a zero in this bit position. Thus, after an interrupt from the timer, the CPU 21 will not accept another such interrupt unless its program causes this reset.

The counter data register 25d is loaded with the contents of the latch 25f via input 25e when either of the following occurs: the GPTSTART bit of the MPCTL register is toggled from 0 to 1 by writing a word to MPCTL by the CPU 21; or the register 25d decrements to zero.

The decrement of register 25d starts when the CPU 21 writes a one to the timer start bit of MPCTL and the previous value of this bit was zero. The local clock at 25b, nominally 3-MHz, is divided by 512 by prescaler 25a and forms a decrement pulse on input 25c of the register 25d. When register 25d decrements to zero, the GPTINT bit 25j is set, and if GPTIEN is set at input 25q, an interrupt is asserted to CPU 21 in line 25r.

On the local bus cycle following the one which decremented register 25d to zero, the register 25d is immediately reloaded with the value in latch 25f. Hence, the period of the general purpose timer interrupts is 512x (cycle time) x (latch 25f value). In a typical situation, this interrupt should be about ten millisecond, i.e. about every thirty thousand machine cycles for the CPU 21. Using registers in RAM 22, the CPU 21 generates a number of different timeouts under its program control, based on this hardware timeout. For example, if no free token or starting delimiter appears on the ring 11 for a given time, then it is assumed that a fault condition has occurred, and fault detector or corrector signals are sent.

35

BUS ARBITRATION

The bus arbiter 26 in the chip 20 is provided to allow the protocol handler chip 16, the system interface chip 30, or another external DMA device to arbitrate for control of the local bus 23. Immediately following system reset by SRESET-, the CPU 21 of chip 20 has control of the bus 23, but the chip 16 and the chip 30 can request and gain control of the bus 23 to perform DMA transfers. Following completion of the DMA cycle or cycles, bus control is returned to the chip 20. In order to maintain maximum utilization of the bandwidth of the local bus 23, the request and grant sequences for future bus cycles are overlapped with the bus cycle currently in progress.

Arbitration and control of the local bus 23 of the arbiter 26 is coordinated by a bus-request/bus-grant handshake using the LBRQ1-, LBGR1, and LBRQ2, LBGR2-handshake pairs on the chip 20, as illustrated in Figures Xa and Xb. The LBRQ1- input from the chip 16 has priority over the LBGR2- input from the system interface chip 30.

The external bus master (chip 16 or chip 30) asserts LBRQ1- or LBRQ2- to request the bus, and the arbiter 26 responds to the bus request by asserting LBGR1- or LBGR2-, then floating the address/data, LEN-, LAL, LI/D-, LRW-lines from the chip 20 at the completion of the bus cycle.

In order of priority, the masters of the local bus 23, as determined by the arbiter 26, are:

- (1) External master on LBRQ1-, the protocol handler chip 16;
- (2) External master on LBRQ2-, the system interface chip 30;
- (3) Internal master from the CPU 21 for microprocessor instruction and data access.

Within the system interface chip 30, the memory-mapped I/O bus controller has priority over the system DNA bus controller.

It is possible that a higher priority bus master may request the local bus 23 while one of the external bus masters has control. That is, the protocol handler chip 16 may request the bus 23 while the system interface chip 30 has control. This could occur when an incoming message is detected at a time when an outgoing message frame is being copied from memory 32 to memory 22 by DMA. In this case, the following occurs:

- 5 (1) the bus arbiter 26 deasserts LBGRI2- in Q3 of the current bus cycle;
- (2) the system interface chip 30 completes the current cycle (possibly with wait states), and tristates its bus signals, leaving LBRQ2- asserted during this time;
- (3) LBGRI1- is asserted;
- 10 (4) when the higher priority master, chip 16, has accomplished its DMA task, such as copying an incoming message frame into the RAM 22, it deasserts LBRQ1-, and the bus arbiter 26 then reasserts LBGRI2-;
- (5) the system interface chip 30 then reassumes control of the bus 23, and completes copying the outgoing data frame.

15

SYSTEM INTERFACE CHIP

20 Referring to Fig. 2, the system interface chip 30 implements two complete bus interfaces, one to the system data bus 34 (and its control bus 35), and the other to the local address/data bus 23. The input and output pins of the chip 30 are described in Table F. Of course, some of the pins are the same as in Table E since some go to the chip 20.

The chip 30 contains two separate controllers. First, a memory-mapped I/O controller with the MMIO register which manages the references by the host system 12 (the host CPU 31) to the memory-mapped registers presented by the adapter 10. The second is a DMA controller which performs DMA transfers between the host system bus 34 and the adapter's local bus 23. With these mechanisms, a variety of different logical interfaces between the host system 12 and the adapter 10 may be implemented with the code executed in local CPU 21 (code stored in ROM 27). The terms memory-mapped I/O and system DMA refer to the data transfer mechanisms as seen from the host system 12 side. This is different from registers in the memory space of local bus 23 on the local side of the chip 30.

For memory-mapped I/O, the adapter 10 appears to the host system 12 as a set of eight consecutive byte addresses. Two registers are dedicated to bit-level status and control information; the program code in ROM 27 executed by the CPU 21 defines the meaning of these bits. The memory-mapped I/O unit 61 also maintains an address register into the local data space of the adapter 10, and the host system 12 may indirectly access any RAM 22 or ROM 27 byte in the memory on the local bus. During system 12 or CPU 31 reads or writes to local memory 22 or the local data space, the CPU 21 is locked out of the local RAM 22. With this approach, command and status task blocks may be written by the host CPU 31 and asynchronously examined by the CPU 21. The memory-mapped I/O unit also places a programmable interrupt vector on the system bus 34 during system bus interrupt acknowledge cycles.

40 The system interface chip 30 also allows for the local CPU 21 to initiate direct memory access between the on-chip RAM 22 and the system memory 32, for incoming frame data or command/status interchange. This DMA is completely under control of the program for CPU 21 stored in the ROM 27.

45 SYSTEM TRANSMIT

To transmit a message, the host system 12 first creates a data frame in its memory 32, under control of the host CPU 31. Then, the CPU 31 writes to the MMIO register of system interface chip 30, which causes an interrupt of the CPU 21, and indicates in the MMIO register the starting address on the system bus 34 of the data frame.

The CPU 21 then sets up a system DMA for system-to-local transfer. This is accomplished by executing program code from the ROM 27 starting at the vector address used for the interrupt. The CPU 21 sets up the following registers: LDMA ADR to contain the system bus 34 address for the data frame; SDMA ADR to contain the local RAM 22 address via local address/data bus 23 for the frame; a SDMA CTL control 55 bit is set to start; SDMA LEN gives the length of the frame. A system DMA through the chip 30 then transfers the data frame from the system RAM memory 32 to the local RAM 22, using a FIFO in the chip 30 so that transfers of data words from the bus 34 to the FIFO are under control of the system control 64 in the chip 30 and synched to the system clock SBCLK, while transfers from the FIFO to the local RAM 22 via bus

23 are under control of the DMA controller in the chip 30 and synched to the local bus clock LBCLK. A cycle on the system control side, if the FIFO is not full, is as follows: the system side controller requests the system bus 34; the host system 12 sends a bus grant to the chip 30; a word (or byte) of data is sent from the system memory 32 into the FIFO; SDMA LEN register is decremented. These cycles continue until the SDMA LEN register is zero, or until the SBERR system bus error pin is asserted, or the DMAHALT bit of the SIFCTL register is set. A cycle on the local control side, if there is data in the FIFO is: the local side controller requests the local bus 23 by a LBRQ1 signal; the LBGR1 is asserted by the chip 20; one word of data is written into the local RAM 22 from the FIFO.

When the entire frame has been copied from the RAM 32 to the RAM 22 in this manner, the system interface chip 30 again interrupts the CPU 21 to indicate that the DMA is complete. The CPU 21 can then operate on the data in local RAM 22 for checking, reformatting, encrypting, etc. When ready to transmit, the CPU 21 sets up the chip 16 for transmit by writing the local starting address and length of the data frame into two registers of the chip 16. The control of transmit now switches to the chip 16. First the transmit FIFO 37 is filled by DMA from chip 16 into RAM 22 of chip 20, using LBRQ2- and LBGR2- controls. The chip 16 then waits for a free token on the input 13 and 13' from the signal line 11. When a free token comes in, the chip 16 changes it to a busy token (sets bit to one), and starts sending data from the FIFO 37 through serializing register 38 and multiplexer 39 to the output 14' and 14. Whenever the FIFO 37 is not full, the chip 16 asserts a request for the local bus by line LBRQ2- and when it gets a LBGR2- from the bus arbiter 26 of the chip 20, a word of data is read from the local RAM 22 to the FIFO 37; the length register is decremented for each word transmitted; this continues until the length register is zero.

The data rate on the ring 11 may be about four 5-bit per second or 0.5 M-byte per second, in a typical example. The transfer rate from the local RAM 22 to the FIFO 37 is at a maximum rate defined by the local bus clock or the cycle time of the CPU 21, which is about 333 nsec, for example, meaning 16-bits are transferred every 333 nsec, which is six M-bytes per sec. Thus, the FIFO 37 can be kept filled during this transfer using only about 0.5/6 or one-twelfth of the cycles of the CPU 21. So the CPU 21 can be doing other operations, merely loosing a cycle now and then for the data transfer to the chip 16 for transmit. Likewise, the host system 12 typically has a system clock of 8-MHz, providing a data transfer rate on the system bus 34 of two M-byte per sec, which is four times that on the signal line 11. Thus, the DMA to transfer the frame from the system RAM 32 to the local RAM 22 is at a higher rate than the off-loading transfer from RAM 22 to the FIFO 37 in chip 16 for transmitting. Accordingly, it is possible to interleave DMA from host system 12 to local RAM 22 with DAM by chip 16 to local RAM 22, although this would not be necessary in most of the transmit operations.

36 TABLE A

THE INSTRUCTION SET

40 DATA TRANSFER INSTRUCTIONS

The MOV instructions are used to transfer data from one part of the memory 2 to another part, or from one location in the memory map to any other. The LOAD instructions are used to initialize registers to desired values. The STORE instructions provide for saving the status register (ST) or the workspace pointer (WP) in a specified workspace register.

Load Immediate - LI

50 Operation: The 16 bit data value located at the address given in the word immediately following the instruction LI is loaded into the workspace register R specified by the 4-bit field of bits 12-15 of this LI instruction.

Applications: The LI instruction is used to initialize a selected workspace register with a program 55 constant such as a counter value or data mask.

Load Interrupt Mask Immediate - LIMI

Operation: The low order 4-bit value (bits 12-15) in the word immediately following this instruction is loaded into the interrupt mask portion of the status register (bits 12-15).

Application: The LIMI instruction is used to initialize the interrupt mask to control which system interrupts will be recognized.

5

Load Workspace Pointer Immediate - LWPI

Operation: The 16-bit value contained in the word immediately following the instruction is loaded into the workspace pointer WP.

Application: LWPI is used to establish the workspace memory area for a section of the program.

MOVE Word - MOV

15

Operation: The word in the source location specified by bits 10-15 of the instruction is transferred to the destination location specified by bits 4-9, without affecting the data stored in the source location. During the transfer, the word (source data) is compared to 0 with the result of the comparison stored in the status register.

Application: MOV is used to transfer data from one part of the system to another part.

MOVE Byte - MOVB

25

Operation: Like MOV except operates on bytes.

Application: MOVB is used to transfer 8-bit bytes from one byte location to another.

Swap Bytes - SWPB

30

Operation: The most significant byte and the least significant byte of the word at the memory location specified by bits 10-15 of the instruction are exchanged.

Application: Used to interchange bytes if needed for subsequent byte operations.

35

Store Status - STST

Operation: The contents of the status register ST are stored in the workspace register specified by bits 12-15 of the instruction.

Application: STST is used to save the contents of status register ST for later reference.

Store Workspace Pointer - STWP

45 Operation: The contents of the workspace pointer WP are stored in the workspace register specified by bits 12-15 of the instruction.

Application: STWP is used to save the contents of the workspace pointer register WP for later reference.

50

ARITHMETIC INSTRUCTIONS

These instructions perform the following basic arithmetic operations: addition (byte or word), subtraction (byte or word), multiplication, division, negation, and absolute value.

55

Add - A

Operation: The data located at the source address specified by bits 10-15 of the instruction is added to the data located at the destination address specified by bits 4-9. The resulting sum is placed in the destination location, and is compared to zero for setting status bits.

Application: Binary addition is the basic arithmetic operation required to generate many mathematical functions. It functions as "pass" when one input is zero. This instruction can be used to develop programs to do multiword addition, decimal addition, code conversion, and so on.

Add Bytes - AB

10 Operation: Like Add but for bytes instead of words. The source byte addressed by bits 12-15 is added to the destination byte addressed by bits 4-9 and the sum byte is placed in the destination byte location. Useful when dealing with subsystems or data that use 8-bit units, such as ASCII codes.

15 Add Immediate - AI

Operation: The 16-bit value contained in the word immediately following the instruction is added to the contents of the workspace register specified by bits 12-15 of the AI instruction.

20 Application: Used to add a constant to a workspace register. Useful for adding a constant displacement to an address contained in the workspace register.

Subtract Words - S

25 Operation: The 16-bit source data (location specified by bits 10-15) is subtracted from the destination data (location specified by bits 4-9) with the result placed in the destination location. The result is compared to 0.

30 Application: Provides 16-bit binary subtraction.

Subtract Bytes - SB

35 Operation: Like S except for bytes instead of words.

Increment - INC

40 Operation: The data located at the source address indicated by bits 10-15 of the INC instruction is incremented and the result is placed in the source location and compared to 0.

45 Application: INC is used to increment byte addresses and to increment byte counters.

Increment by Two - INCT

50 Operation: Two is added to the data at the location specified by the source address in bits 10-15 and the result is stored at the same source location and is compared to 0 to set status ST. Similar to INC.

Application: This can be used to increment word addresses, though autoincrementing on word instructions does this automatically.

Decrement - DEC

55 Operation: One is subtracted from the data at the location specified by bits 10-15, the result is stored at that location and is compared to 0. Similar to INC.

Application: Most often used to decrement byte counters or to work through byte addresses in descending order.

Decrement by Two - DECT

Operation: Two is subtracted from the data at the location specified by bits 10-15 and the result is stored at that location and is compared to 0. Similar to INC.

6 Application: This instruction is used to decrement word counters and to work through word addresses in descending order.

Negate - NEG

10 Operation: The data at the address specified by bits 10-15 of the instruction is replaced by its two's complement, and the result is compared to zero for setting status ST.

Application: NEG is used to form the two's complement of 16-bit numbers.

15 Absolute Value - ABS

Operation: The data at the address specified by bits 10-15 of the instruction is compared to zero to set status ST. Then the absolute value of this data is placed in the same location.

20 Application: Used to test the data in the specified location and then replace the data by its absolute value. This could be used for unsigned arithmetic algorithms such as multiplication.

Multiply - MPY

25 Operation: The 16-bit data at the source address designated by bits 10-15 is multiplied by the 16-bit data contained in the destination workspace register R specified by bits 6-9. The unsigned binary product (32-bits) is placed in workspace registers R and R + 1.

30 Application: MPY can be used to perform 16-bit by 16-bit binary multiplication. Several such 32-bit subproducts can be combined in such a way to perform multiplication involving larger multipliers and multiplicands such as a 32-bit by 32-bit multiplication.

Divide - DIV

35 Operation: The 32-bit number contained in workspace registers R and R + 1 (where R is a destination address specified at bits 6-9) is divided by the 16-bit divisor contained at the source address specified by bits 10-15. The workspace register R then contains the quotient and workspace R + 1 contains the 16-bit remainder. The division will occur only if the divisor at the source address is greater than the data contained in R.

40 Application: DIV provides basic binary division of a 32-bit number by a 16-bit number.

Signed Multiply - MPYS.

45 Operation: The signed two's complement integer in workspace register R0 is multiplied by the signed two's complement integer specified by the source address (bits 10-15). The result is a signed 32-bit product which is placed in workspace register R0 (for the 16 MSB's) and workspace register R1 (the 16 LSB's). The result is compared to zero.

50 Application: Provides signed multiplication for the system.

Signed Divide - DIVS

55 Operation: The signed 32-bit two's complement integer (dividend) in workspace registers 0 and 1 is divided by the signed 16-bit two's complement integer (divisor) specified by the source address (bits 10-15). The signed quotient is placed in workspace register R0 and the signed remainder is placed in workspace register R1.

COMPARISON INSTRUCTIONS

These instructions are used to test words or bytes by comparing them with a reference constant or with another word or byte. Such operations are used in certain types of division algorithms, number conversion, and in recognition of input command or limit conditions.

Compare Words - C

10 Operation: The two's complement 16-bit data specified by the source address at bits 10-15 is compared to the two's complement 16-bit data specified by the destination at bits 4-9, and appropriate status ST bits are set based upon the results. The contents of both locations remain unchanged.

Application: The need to compare two words occurs in such system functions as division, number conversion, and pattern recognition.

15

Compare Bytes - CB

20 Operation: The two's complement 8-bit byte at the source address (bits 10-15) is compared to the two's complement 8-bit byte at the destination address (bits 4-9), and status ST is set.

Compare Immediate - CI

25 Operation: CI compares the contents of the specified workspace register R (defined by bits 12-15) to the value contained in the word immediately following the instruction, and sets status bits accordingly.

Application: CI is used to test data to see if system or program limits have been met or exceeded or to recognize command words.

30

Compare Ones Corresponding - COC

Operation: The data in the source location addressed by bits 10-15 acts as a mask for the bits to be tested in workspace register R specified by bits 6-9. That is, only the bit position that contain ones in the source data will be checked in R. Then, if R contains ones in all the bit positions selected by the source data, the equal (EQ) status bit will be set to 1.

Application: COC is used to selectively test groups of bits to check the status of certain sub-systems or to examine certain aspects of data words.

40

Compare Zeroes Corresponding - CZC

Operation: The data located in the source address specified by bits 10-15 act as a mask for the bits to be tested in the workspace register R specified by bits 6-9. That is, only the bit positions that contain ones in the source data are the bit positions to be checked in R. Then if R contains zeroes in all the selected bit positions, the equal (EQ) status bit will be set to 1.

Application: Similar to the COC instruction.

50

LOGIC INSTRUCTIONS

The Logic Instructions allow the processor to perform boolean logic for the system. Since AND, OR, INVERT, and Exclusive OR (XOR) are available, any boolean functions can be performed on system data.

55

AND Immediate - ANDI

Operation: The bits of the workspace register specified by bits 12-15 are logically ANDed with the

corresponding bits of the 16-bit binary constant value contained in the word immediately following the instruction. The 16-bit result is compared to zero and is placed in the register R.

Application: ANDI is used to zero all bits that are not of interest and leave the selected bits (those with ones in value) unchanged. This can be used to test single bits or isolate portions of the word, such as a 6 four-bit group.

OR Immediate - ORI

10 Operation: The bits of the specified workspace register R are ORed with the corresponding bits of the 16-bit binary constant contained in the word immediately following instruction. The 16-bit result is placed in R and is compared to zero.

Application: Used to implement the OR logic in the system.

15 **Exclusive OR - XOR**

Operation: The exclusive-OR function is performed between corresponding bits of the source data addressed by bits 10-15 and the contents of workspace register specified by bits 6-9. The result is placed 20 in workspace register R and is compared to zero to set status ST.

Invert - INV

25 Operation: The bits of the data addressed by the source address bits 10-15 are replaced by their complement. The result is compared to zero and is stored at the source location.

Application: INV is used to form the one's complement of 16-bit binary numbers, or to invert system data.

30 **Clear - CLR**

Operation: Zeroes are placed in the memory location specified by bits 10-15 of the instruction.

Application: CLR is used to set problem arguments to zero and to initialize memory locations to zero 35 during system starting operations.

Set to One - SETO

40 Operation: All ones (or hex FFFF) are placed in the memory location specified by bits 10-15 of the instruction.

Set Ones Corresponding - SOC

45 Operation: This instruction performs the OR operation between corresponding bits of the source data address defined by bits 10-15 and the destination data addressed defined by bits 4-9. The result is compared to zero and is placed in the destination location.

Application: Provides the OR function between any two words in memory.

50

Set Ones Corresponding Byte - SOCB

55 Operation: Like SOC except used for bytes instead of words. The logical OR is performed between corresponding bits of the byte addressed by source address bits 10-15 and the byte addressed by the destination address of bits 4-9 with the result compared to zero and placed in destination location.

Set to Zeroes Corresponding - SZC

Operation: The source data addressed by bits 10-15 forms a mask for this operation. The bits in the destination data (addressed by bits 4-9) that correspond to the one bits of the source data (addressed by bits 10-15) are cleared. The result is compared to zero and is stored in the destination location.

5 Application: SZC allows the programmer to selectively clear bits of data words. For example, when an interrupt has been serviced, the interrupt request bit can be cleared by using the SCZ instruction.

10 Set to Zeroes Corresponding, Bytes - SZCB

Operation: Like SZC except for bytes instead of words. The byte addressed by the source address bits 10-15 will provide a mask for clearing certain bits of the byte addressed by the destination address bits 4-9. The bits in the destination byte that will be cleared are the bits that are one in the source byte. The result is compared to zero and is placed in the destination byte.

15 These instructions are used to perform simple binary multiplication and division on words in memory and to rearrange the location of bits in the word in order to examine a given bit with the carry (C) status bit.

20 Shift Right Arithmetic - SRA

Operation: The contents of the workspace register R specified by bits 12-15 are shifted right by a number of times specified by bits 8-11, referred to as Cnt, filling the vacated bit position with the sign (most significant bit) bit. The shifted number is compared to zero and status bits set.

25 Number of Shifts = Cnt (number from 0 to 15 contained in the instruction) specifies the number of bits shifted, unless Cnt is zero in which case the shift count is taken from the four least significant bits of workspace register 0. If both Cnt and these four bits are 0, a 16-bit position shift is performed:

Application: SRA provides binary division by a power of two defined by Cnt.

30 Shift Left Arithmetic - SLA

Operation: The contents of workspace register R specified by bits 12-15 are shifted left Dnt times (or if Dnt = 0, the number of times specified by the least four bits of R0) filling the vacated positions with zeroes.

35 Dnt is specified by bits 8-11. The carry contains the value of the last bit shifted out to the left and the shifted number is compared to zero and status bits set.

Application: SLA performs binary multiplication by a power of two defined by Cnt.

40 Shift Right Logical - SRL

Operation: The contents of the workspace register R specified by bits 12-15 are shifted right Cnt times, where Cnt is specified by bits 8-11, (or if Cnt = 0, the number of times specified by the least four bits of R0) filling in the vacated positions with zeroes. The carry contains the value of the last bit shifted out the right

45 and the shifted number if compared to zero.

Application: Performs binary division by a power of two defined by Cnt.

Shift Right Circular - SRC

50 Operation: Workspace register R defined by bits 12-15 is right shifted Cnt time defined by bits 8-11. On each shift the bit shifted and the shifted number is compared to 0. The number of shifts to be performed is the number Cnt, or if Cnt = 0, the number contained in the least significant four bits of R0.

55 Application: SRC can be used to examine a certain bit in the data word, change the location of 4-bit groups, or swap bytes.

UNCONDITIONAL BRANCH INSTRUCTIONS

These instructions give the programmer the capability of choosing to perform the next instruction in sequence or to go to some other part of the memory to get the next instruction to be executed. The branch can be a subroutine type of branch, in which case the programmer can return to the point from which the branch occurred.

5

Branch - B

Operation: The source address, defined by bits 10-15, is placed in the program counter, causing the next instruction to be obtained from the location specified by this source address.

Application: This instruction is used to jump to another part of the program when the current task has been completed.

16 Branch and Link - BL

Operation: The source address defined at bits 10-15 is placed in the program counter and the address of the instruction following the BL instruction is saved in workspace register R11.

Application: This is a shared workspace subroutine jump. Both the main program and the subroutine use the same workspace registers. To get back to the main program at the branch point, a BL instruction can be used at the end of the subroutine which causes the R11 contents (old PC value) to be loaded into the program counter.

25 Unconditional Jump - JMP

Operation: The signed displacement defined by bits 8-15 is added to the current contents of the program counter PC to generate the new value of the program counter. The location jumped to must be within -128 to +127 words of the present location.

Application: If the subprogram to be jumped to is within 128 words of the JMP instruction location, the unconditional JMP is preferred over the unconditional branch since only one memory word (and one memory reference) is required for the JMP while two memory words and two memory cycles are required for the B instruction. Thus, the JMP instruction can be implemented faster and with less memory cost than can the B instruction.

35

Execute - X

Operation: The instruction located at the source address specified by bits 10-15 is executed.

Application: X is useful when the instruction to be executed is dependent on a variable factor.

Branch and Load Workspace Pointer - BLWP

45 Operation: The word specified by the source address bits 10-15 is loaded into the workspace pointer WP and the next word in memory (source address + 2) is loaded into the program counter PC to cause the branch. The old workspace pointer is stored in the new workspace register R13, the old PC value is stored in the new workspace register R14, and the status register is stored in new workspace register R15.

Application: This is a context switch subroutine jump with the transfer vector location specified by the source address. It uses a new workspace to save the old values of WP, PC, and ST (in the last three registers). The advantage of this subroutine jump over the BL jump is that the subroutine gets its own workspace and the main program workspace contents are not distributed by subroutine operations.

Operation: Bits 6-9 specify which extended operation transfer vector is to be used in the context switch branch from XOP to the corresponding subprogram. The effective source address, bits 10-15, is placed in R11 of the subprogram workspace in order to pass an augment or data location to the subprogram.

Application: This can be used to define a subprogram that can be called by a single instruction. As a result, the programmer can define special purpose instructions to augment the standard instruction set for the processor.

Return with Workspace Pointer - RTWP

Operation: This is a return from a context switch subroutine. It occurs by restoring the WP, PC, and ST register contents by transferring the contents of subroutine workspace registers R13, R14 and R15, into the WP, PC and ST registers, respectively.

Application: This is used to return from subprograms that were reached by a transfer vector operation such as an interrupt, extended operation, or BLWP instruction.

70 CONDITIONAL JUMP INSTRUCTIONS - JH, JL, JEE, JLE, JGT, JLT, JEQ, JNE, JOC, JNC, JNO, JOP

These instructions perform a branching operation to a location defined by bits 8-15 only if certain status bits meet the conditions required by the jump. These instructions allow decision making to be incorporated into the program. The conditional jump instruction mnemonics are summarized below along with the status bit conditions that are tested by these instructions.

Operation: If the condition indicated by the branch mnemonic (specified by bits 4-7) is true, the jump will occur using relative addressing as was used in the unconditional JMP instruction. That is, the bits 8-15 define a displacement that in added to the current value of the program counter to determine the location of the next instruction, which must be within 128 words of the jump instruction.

20

Status Bits Tested by Instructions

25	Mnemonic	L	A	EQ	C	OV	OP	Jump if:	CODE*
	JH	X	-	X	-	-	-	L EQ=1	B
30	JL	X	-	X	-	-	-	L +EQ=0	A
	JHE	X	-	X	-	-	-	L +EQ=1	4
	JLE	X	-	X	-	-	-	L +EQ=1	2
	JGT	-	X	-	-	-	-	A =1	5
35	JLT	-	X	X	-	-	-	A +EQ=0	1
	JEQ	-	-	X	-	-	-	EQ=1	3
	JNE	-	-	X	-	-	-	EQ=0	6
40	JOC	-	-	-	X	-	-	C=1	8
	JNC	-	-	-	X	-	-	C=0	7
	JNO	-	-	-	-	X	-	OV=0	9
45	JOP	-	-	-	-	-	X	OP=1	C

*CODE is the CODE field (bits 4-7) of the OPCODE to

50 generate the machine code for the instruction.

Application: Most algorithms and programs with loop counters require these instructions to decide which sequence of instructions to do next.

55

TABLE B

	INSTRUCTION	OPCODE	INSTRUCTION	OPCODE
5	SOCB	1111XXXXXXXXXXXX	SRA	00001000XXXXXXX
	SOC	1110XXXXXXXXXXXX	ABS	0000011101XXXXXX
	MOVB	1101XXXXXXXXXXXX	SETO	0000011100XXXXXX
10	MOV	1100XXXXXXXXXXXX	SWPB	0000011011XXXXXX
	AB	1011XXXXXXXXXXXX	BL	0000011010XXXXXX
	A	1010XXXXXXXXXXXX	DECT	0000011001XXXXXX
	CB	1001XXXXXXXXXXXX	DEC	0000011000XXXXXX
	C	1000XXXXXXXXXXXX	INCT	0000010111XXXXXX
15	SB	0111XXXXXXXXXXXX	INC	0000010110XXXXXX
	S	0110XXXXXXXXXXXX	INV	0000010101XXXXXX
	SZCB	0101XXXXXXXXXXXX	NEG	0000010100XXXXXX
	SZC	0100XXXXXXXXXXXX	CLR	0000010011XXXXXX
	DIV	001111XXXXXXXXXX	X	0000010010XXXXXX
20	MPY	001110XXXXXXXXXX	B	0000010001XXXXXX
	XOP	001011XXXXXXXXXX	BLWP	0000010000XXXXXX
	XOR	001010XXXXXXXXXX	RTWP	00000011100XXXXXX
	CZC	001000XXXXXXXXXX	RSET	00000011011XXXXXX
	COC	001000XXXXXXXXXX	IDLE	00000011010XXXXXX
25	JOP	00011100XXXXXXXX	LIMI	00000011000XXXXXX
	JH	00011011XXXXXXXX	LWPI	0000001011100000
	JL	00011010XXXXXXXX	STST	00000010110XXXXXX
	JNO	00011001XXXXXXXX	STWP	00000010101XXXXXX
	JOC	00011000XXXXXXXX	CI	00000010100XXXXXX
30	JNC	00010111XXXXXXXX	ORI	00000010011XXXXXX
	JNE	00010110XXXXXXXX	ANDI	00000010010XXXXXX
	JGT	00101011XXXXXXXX	AI	00000010001XXXXXX
	JHE	00010100XXXXXXXX	LI	00000010000XXXXXX
	JEQ	00010011XXXXXXXX	MFYS	0000000111XXXXXX
35	JLE	00010010XXXXXXXX	DIVS	0000000110XXXXXX
	JLT	00010001XXXXXXXX		
	JMP	00010000XXXXXXXX		
	SRC	00001011XXXXXXXX		
	SLA	00001010XXXXXXXX		

40

45

50

55

TABLE C
ALU FUNCTIONS

CONTROL LINES TO ALU				ALU FUNCTION		
	ALU 1	ALU 2	ALU 3	ALU 4	H3 Logic = 1	H3 Logic = 0
10	0	0	0	0	1	
	0	0	0	1	$A + B$	
	0	0	1	0	$A + \bar{B}$	
15	0	0	1	1	A	
	0	1	0	0	$\bar{A} + B$	
	0	1	0	1	B	
20	0	1	1	0	$\bar{A} \oplus B$	ADD
	0	1	1	1	$A \cdot B$	
	1	0	0	0	$\bar{A} + \bar{B}$	
25	1	0	0	1	$A \oplus B$	$A \text{ sub } B$
	1	0	1	0	\bar{B}	
	1	0	1	1	$A \cdot \bar{B}$	
30	1	1	0	0	\bar{A}	
	1	1	0	1	$\bar{A} \cdot B$	
	1	1	1	0	$\bar{A} + \bar{B}$	
35	1	1	1	1	0	

40

45

50

55

50
45
40
35
30
25
20
15
10
5TABLE D
TRUTH TABLE FOR TEST MODE

PATTERN	CONTROLS		DATA		MEMORY OUT		FAULTED = $M(T+1)$	
	<u>IN-A</u>	<u>IN-Z</u>	<u>IN-A</u>	<u>IN-B</u>	<u>STI</u>	<u>M(T+1)</u>	<u>IN-A</u>	<u>STI</u>
T	0	0	0	0	1	1	0	0
U	1	0	0	0	1	0	1	1
V	1	0	1	0	1	1	0	1
W	1	0	0	0	1	0	0	0
X	0	1	0	0	1	0	1	1
Y	0	1	0	1	1	1	1	0
Z	0	1	0	0	1	0	0	1

55

TABLE E
PIN DEFINITIONS FOR MESSAGE PROCESSOR CHIP

5	PIN	FUNCTION
10	LAL	Address Latch for Local Bus. When LAL goes active (high) the address appearing on the address/data bus is latched. LAL can originate either within the chip 20 or from external to the chip 20.
15	LEN-	Data Enable for Local Bus. When LEN- goes active-low, the data is valid on the local address/data bus 23.
20	LI/D-	Instruction or Data on Local Bus. When high, LI/D- indicates that an instruction fetch is in progress. When low, a data fetch is in progress. Can be used to select between ROM 27 and RAM 22.
25	LR/W-	Read or Write for Local Bus. When high, LR/W- indicates that a read cycle is being implemented. When low, a write cycle is being implemented.
30	LNMI-	Non-Maskable Interrupt for Local Bus. When this goes active-low, the CPU 21 executes a non-maskable interrupt, jumping to a vector address.
35	LIRQ	Interrupt Request for Local Bus. Three pins. Contain the incoming interrupt code to define the interrupt level.
40	LBRDY	Local Bus Ready. External devices introduce wait states by holding LBRDY low. The CPU 21 continues wait states until LBRDY goes high.
45	LBRQ1- LBRQ2-	Local Bus Request 1 and 2. These inputs to chip 20 are driven active-low by the protocol handler chip 16 and the system interface chip 30, respectively, to request control of the local bus 23. LBRQ1- has priority over LBRQ2-. The local bus arbiter 26 of the chip 20 samples both signals at a given clock phase, and asserts either LBGR1- or LBGR2- on the next phase.
50	LBGR1- LBGR2	Local Bus Grant 1 and 2. These outputs from the chip 20 are driven active-low by the local bus arbiter 26 of the chip 20 in response to a bus request LBRQ1- or LBRQ2-, respectively. It indicates that the requesting device 16 or 20 may use the following cycle, if LBRDY is also asserted (high).

	LBSYNC-	Local Bus Synchronization. Used only for test of chip 20. When active-low this pin forces the CPU 21 clock generator into the Q4 state.			
5	TEST-	For normal operation, TEST- and TEST 0 are held high. For test operations, the input test codes shown in the following list are implemented.			
10	<u>TEST-</u>	<u>TEST0</u>	<u>TEST1</u>	<u>TEST2</u>	
15	1	1	X	X	Normal Operation
20	0	0	0	0	Module-in-place test mode. All output pins of the chip 20 are floated.
25	0	0	0	1	Entry point test mode. When MPRSET- goes active then is deactivated, if this code is present an external tester has control of the address input 46 to the control ROM 45 of the local CPU 21.
30	0	0	1	0	Test mode for general purpose timer 25 in chip 20. A microcode sequence is executed for testing the timer.
35	0	0	1	1	Not specified.
40	0	1	0	X	Dump for Control ROM 45. When MPRSET- is deasserted, the CPU 21 enters a state that dumps the contents of the control ROM 45, shifting each 130-bit control word out serially on a pin (such as LBGR2- pin). The TEST2 pin is toggled to increment the control ROM address register.
45	0	1	1	X	Refresh test mode for RAM 22. When MPRSET- is deasserted, the chip 20 enters a state that permits testing of the internal nodes in the RAM 22 refresh logic. TEST2 provides timing information.

TABLE F
PIN DEFINITION FOR SYSTEM INTERFACE CHIP

PIN	FUNCTION
5 SI/M-	System Mode select. This input pin, if held high, at I-mode, causes the chip 30 to be compatible with an 8086 or 8088 microprocessor chip 31. If held low, at M-mode, SI/M- causes the chip 30 to operate in a format compatible with a 68000 microprocessor chip 31.
10 S5/16-	System 8/16-bit bus select. If held low, an interface mode with a 16-bit data bus 34 is selected (the 8086 device). If held high, an 8-bit data bus 34 is selected (for 8088 device, for example).
15 SRSET	System Reset. This input pin places the entire adapter 10 in a known initial state. The chip 30 passes this signal through to the chip 20 via the LRESET- pin.
20 SCS-	System Chip Select. An input pin to chip 30 from the host CPU 31 functioning as chip select to allow the host processor 12 to execute a memory-mapped I/O to the chip 30 for read or write.
25 SRS0	System Register Select. These three inputs to the chip 30 select the word or byte to be addressed during a memory-mapped I/O access from host CPU 31 to the chip 30.
30 SRS1	
35 SRS2	
SBHE-/SRNW	System Byte High Enable or Read/Not Write. In the I-mode (SI/M- = 1) this pin serves as an active-low byte-high-enable signal SBHE-. In the M-mode (SI/M- = 0), serves as a control signal which is high to indicate a read cycle and low to indicate a write cycle. The chip 30 drives this pin as an output during DMA onto the system bus 34; it is an input during memory-mapped I/O cycles from host CPU 31 to chip 30.
40 SWR-/SLDS	System Write Strobe or Lower Data Strobe. In I mode, this pin serves as the active-low write strobe. In M-mode, serves as the active-low lower data strobe. An input to chip 30 during memory-mapped I/O, and an output during DMA.
45 SRD-/SUDS	System Read Strobe or Upper Data Strobe indicator. In the I-mode, this pin is an active-low strobe indicating that a read cycle is performed on the system bus 34. In M-mode, an active-low strobe indicating that data is transferred on the most significant byte of the
50	

system bus 34. An input to chip 30 during memory-mapped I/O, and an output during DMA.

5 SRAS/SAS- System Register Address Strobe or Memory Address Strobe. In I-mode, this pin serves as the system register address strobe, by means of which SCS-, SRS0 to SRS2, and SBHE- are latched. In I-mode, the SBHE- input is also latched. In a minimum-chip system, SRAS is typically tied to the ALE output of the host CPU 31. This latching capability is easily defeated, as is usually desired in expanded 8086/8088 systems supporting non-multiplexed address and data busses. The internal latch for these inputs remains transparent as long as SRAS remains high, permitting this pin to be strapped high and the signals at the SCS-, SBHE-, and SRS0 to SRS2 inputs to be applied independent of an ALE strobe from the host CPU 31. In M-mode, this pin is an active-low address strobe, which is an input during memory-mapped I/O and an output during DMA.

10 SRDY-/ SDTACK- System Bus Ready or Data Transfer Acknowledge. In I-mode, this pin serves as an active-low bus ready signal. In M-mode, this pin serves as the active-low data transfer acknowledge signal. The purpose of the SRDY- and SDTACK- signals is to indicate to the bus master that a data transfer is complete. SRDY-/SDTACK- is internally synchronized to SBCLK; it must be asserted before the falling edge of SBCLK in state T2 in order to prevent a wait state. SRDY-/SDTACK- is an output when the system interface chip 30 is selected for memory-mapped I/O, and an input otherwise.

15 SALE System Address Latch Enable. At the start of each DMA cycle, this output from chip 30 provides the enable pulse used to externally latch the sixteen LSBs of the address from the multiplexed address/data lines.

20 SXAL System Extended Address Latch. This output provides the enable pulse used to latch (external to the chips 30) the eight address extension bits of the 24-bit system address during DMA. SXAL is activated prior to the first cycle of each block DMA transfer, and thereafter as necessary (whenever an increment of the DAM address counter causes a carry-out of the lower 16-bits).

25 SDIR System Data Direction. This output provides to the external data buffers a signal indicating the direction in which the data is moving. During memory-mapped I/O writes and DMA reads,

5 **SDDIR** **DIRECTION** **MMIO** **DMA**
 10 0 input write read
 1 output read write

When the system interface chip 30 is not involved in a memory-mapped I/O or DMA operation, then SDDIR is high by default.

5	SDBEN-	System Data Bus enable. This output provides to data buffers in system 12 external to the chip 30, the active-low enable signal that causes them to leave the high-impedance state and begin transmitting data. This output is activated during both memory-mapped I/O and DMA.
10	SOWN-	System Bus Owned. This output goes active-low during DMA cycles to indicate to external devices that the chip 30 has control of the system bus. SOWN- drives the enable signal of the bus transceivers chips which drive the address and bus control signals.
15	SBCLK	System Bus Clock. This is the external clock signal with which the chip 30 synchronizes its bus timing for both memory-mapped I/O and DMA transfers.
20	SHRQ/ SBRQ-	System Hold Request or Bus Request. This output is used to request control of the system bus 34 in preparation for a DMA transfer. In I-mode, it is an active-high hold request, as defined in the standard 8086/8088 interface. In M-mode, it is an active-low bus request, as defined in the standard 68000 interface.
25	SHLDA/ SBGR-	System Hold Acknowledge or Bus Grant. In I-mode, this active-high pin indicates that the DMA hold request has been acknowledged, in accordance with the standard 8086/8088 interface. In M-mode, this pin is an active-low bus grant, as defined in the standard 68000 interface. In either mode, it is internally synchronized to SBCLK.
30	SBBSY-	System Bus Busy. This input signal samples the value of the 68000-style Bus Grant Acknowledge (BGACK-) signal. The chip 30 must sample SBBSY- high before it drives the system bus 34. Its operation is defined for both I-mode and M-mode.
35	SBRLS-	System Bus Release. This input is driven active-low during DMA to indicate that a higher-priority device requires the system bus
40		
45		
50		

34 and the chip 30 to release the bus as soon as possible. When the chip 30 is not performing DMA, this input is ignored. Its operation is defined for both I- and M-modes synchronized to SCLK.

5 SINTR/ System Interrupt Request. The chip 30 activates SIRQ- this output to signal an interrupt request to the host processor 12. In I-mode, this pin is active-high; it is active-low in M-mode.

10 SIACK- System Interrupt Acknowledge. This input is driven active-low by the host processor 12 to acknowledge the interrupt request from the chip 30. The chip 30 responds to this signal by gating its interrupt vector onto the system bus 34. System busses not requiring an interrupt cycle may strap SIACK- high.

15 SBERR- Bus Error. This input is driven active-low during a DMA cycle to indicate to the chip 30 that the cycle must be terminated abnormally. It corresponds to the Bus Error signal of the 68000 microprocessor. It is internally synchronized to SCLK. It is sampled in both I- and M-modes.

20 SADH' System Address/Data Bus (high byte). This is the 0 to 7 most significant byte of the 16-bit address/data bus 34. In I-mode, it is attached to the host system address/data bus 34 bits 15-8. In M-mode, it is attached to the host system address/data bus 34 bits 0-7 (using 68000 standard bit numbering conventions).

25 SADL System Address/Data Bus (low byte). This is the 0 to 7 least significant byte of the 16-bit address/data bus. In I-mode, it is attached to the host system address/data bits 7-0 (using 8086 standard bit numbering conventions.)

30 SPE System Parity High. Contains an odd-parity bit for each data or address byte transmitted over SADH 0 to 7.

35 SPL System Parity Low. Contains an odd-parity bit for each data or address byte transmitted over SADL 0 to 7.

40 LBCLK1 Local Bus Clock 1 and Local Bus Clock 2. These LBCLK2 signals are the input clock for all local bus transfers. LBCLK1 follows LBCLK2 by 90 degrees.

45 LAL Local Address Latch Enable. At the beginning of each local bus cycle, LAL is driven high and then low to strobe the address local bus 23 on and into an external latch. The LAL signal is an output of the chip 30 when the chip 30 controls

50

the local bus 23 and input to the chip 30 otherwise.

55

Claims

1. A digital data communications system comprising:
 a signal path (11),
 5 a plurality of stations (10) connected to said signal path, each station receiving digital data from the path and transmitting digital data to said path, each said station comprising
 a main processor (12) containing a main CPU (31) and a main memory (32), and having system bus means (34) interconnecting said CPU and said main memory,
 a communications processor (20) containing a local CPU (21) and local memory means (22,27), and
 10 having local bus means (23,24) interconnecting said local CPU and said local memory means,
 first memory access means (30) connecting said system bus means to said local bus means whereby said main CPU can transfer data to said local memory means and said local CPU can transfer data to said main memory, and
 second memory access means (16) connected to said local bus means to access said local memory
 15 means for data transfer operations;
 characterized in that said communications processor further includes a timer (25) having count register means (25d) connected to said local bus means and having a means (25k,25p) to interrupt said local CPU when said count register means (25d) reaches a predetermined value, and
 bus arbiter means (26) controlling access to said local bus means,
 20 said first and second memory access means being operably coupled to said bus arbiter means, and
 said bus arbiter means receiving bus request signals individually from said first memory access means, said second memory access means, and said local CPU and granting access to said local bus means in a predetermined order of priority.
- 25 2. A digital data communications system as set forth in claim 1, further characterized in that said bus arbiter means in response to receiving bus request signals individually from said first memory access means, said second memory access means and said local CPU grants access to said local bus means in that order of priority of said second memory access means having the highest priority, said first memory access means having the next higher priority, and said local CPU having the lowest order of priority.
3. A digital data communications system as set forth in either of claims 1 or 2, further characterized in that said local memory means of said communications processor includes a local read/write memory (22), said local memory means further including a local read-only-memory (270, and
 35 said first memory access means connecting said system bus means to said local bus means whereby said main CPU can transfer data to said local read/write memory and said local CPU can transfer data to said main memory.
4. A digital data communications system as set forth in claim 1, further characterized by receiving means (15) having an input (13) coupled to said signal path and a receive buffer memory (18), said receive buffer memory being coupled to the input of said receiving means for receiving incoming data from said signal path,
 transmitting means (15) having an output (14), coupled to said signal path, and a transmit buffer memory (37) for receiving data at an input thereof for transmission to the output (14) of said
 45 transmitting means, and
 said communications processor including transmit/ receive control means responsive to receipt of data from said signal path in said receiving means,
 said second memory access means connecting said receive buffer memory and said transmit buffer memory to said local bus means and responsive to said transmit/ receive control means to access said local read/write memory for writing received data from the receive buffer memory to the local read/write memory and reading transmit data from the local read/write memory to the transmit buffer memory.
5. A digital data communications system as set forth in claim 4, further characterized in that a frame of a plurality of bytes of data to be transmitted is created in said main memory under control of said main CPU, and such frame is transferred upon initiation by said main CPU via said system bus means and said local bus means into said local read/write memory, and subsequently such frame is transferred under control of said transmit/ receive control means via said local bus means and said transmit buffer memory to said transmitting means.
 55

6. A digital data communications system as set forth in claim 5, further characterized in that said frames of data created in said main memory and transmitted on said signal path by said stations include a destination address identifying a particular one of the stations, and a source address identifying the originating station.
- 5 7. A digital data communications system as set forth in any of claims 4 to 6, further characterized in that data is transmitted on said signal path in a bit-serial format, said receiving means including serial-to-parallel conversion means (17), said transmitting means including parallel-to-serial conversion means (38), and said system bus means and said local bus means being parallel buses.
- 10 8. A digital data communications system as set forth in any of claims 3 to 7, further characterized in that said local read-only-memory is coupled to said local bus means with said local CPU executing instructions fetched from said local read-only-memory via said local bus means as controlled by said bus arbiter means, and said main CPU executing instructions fetched from said main memory via said system bus means.
- 15 9. A digital data communications system as set forth in any of claims 3 to 8, further characterized in that said first memory access means includes memory-mapped input/ output means and a control register, said main CPU accessing said control register via said memory mapped input/output means, and direct memory access of said local read/write memory from said system bus means being initiated by said control register means in response to the writing of data thereto by said main CPU.
- 20 10. A digital data communications system as set forth in any of claims 3 to 9, further characterized in that said first memory access means comprises a system interface device (30) coupling said local bus means to said system bus means for transferring data between said main memory and said local memory means, and said second memory access means comprising a data communications interface device (16) coupling said local bus means to said signal path for transferring data between said signal path and said local memory means.
- 25 11. A digital data communications system as set forth in claim 10, further characterized in that said local read-only-memory is disposed within said data communications interface device, and said system interface device including a memory-mapped input-output channel by which said main CPU accesses said local memory means to transfer command and status information, and a direct memory access channel by which data is transferred between the main memory and the local memory means under control of the system interface device after initiation of the transfer by the main CPU or the local CPU, and said data communications interface device including a direct memory access channel by which data is transferred between said signal path and said local memory means under control of the data communications interface device after initiation of the transfer by the local CPU or the data communications interface device.
- 30 12. A digital data communications system as set forth in either of claims 10 or 11, further characterized in that said local bus means includes a local multiplexed address/data bus (19,28) connecting said local CPU and said local memory means, said local memory means including said read-only-memory for storing program instructions executed by said local CPU and said read/write memory for temporarily storing message data, said system interface device coupling said local address/data bus to said system bus means for transferring message data between said main memory and said local memory means by direct memory access in both directions, and said data communications interface device coupling said local bus means to said signal path for transferring data between said signal path and said local memory means by direct memory access in both directions.
- 35 13. A digital data communications system as set forth in any preceding claim, further characterized in that said signal path is a closed, one-way path, and each station receives data from an upstream direction in the signal path and transmits data to said signal path in a downstream direction, access to said signal path being controlled by a token included in the transmitted data and passed from

station to station around the closed path.

14. A digital data communications system as set forth in any preceding claim, further characterized in that said communications processor comprises a single-chip microprocessor device.

5

Revendications

1. Système de transmission de données numériques comportant :
 - 10 un circuit de signaux (11), plusieurs postes (10) connectés audit circuit de signaux, chaque poste recevant des données numériques provenant du circuit et émettant des données numériques sur ledit circuit, chacun desdits postes comportant :
 - 15 un processeur principal (12) qui contient une CPU principale (31) et une mémoire principale (32) et comprenant des moyens de lignes omnibus de système (34) interconnectant ladite CPU et ladite mémoire principale,
 - 20 un processeur de transmission (20) qui contient une CPU locale (21) et des moyens de mémoire locale (22, 27) et comprenant des moyens de lignes omnibus locales (23, 24) interconnectant ladite CPU locale et lesdits moyens de mémoire locale,
 - 25 des premiers moyens d'accès à la mémoire (30) connectant lesdits moyens de lignes omnibus de système auxdits moyens de lignes omnibus locales de manière que ladite CPU principale puisse transférer des données vers lesdits moyens de mémoire locale et que ladite CPU locale puisse transférer des données vers ladite mémoire principale, et
 - 30 des seconds moyens d'accès à la mémoire (16) connectés auxdits moyens de lignes omnibus locales pour accéder auxdits moyens de mémoire locale pour les opérations de transfert de données ; caractérisé en ce que ledit processeur de transmission comporte en outre un temporisateur (25) comportant des moyens de registre de comptage (25d), connectés auxdits moyens de lignes omnibus locales et comprenant des moyens (25k, 25p) pour interrompre ladite CPU locale quand lesdits moyens de registre de comptage (25d) atteignent une valeur prédéterminée, et
 - 35 des moyens d'arbitrage de lignes omnibus (26) commandant l'accès auxdits moyens de lignes omnibus locale, lesdits premiers et lesdits seconds moyens d'accès à la mémoire étant couplés en fonctionnement avec lesdits moyens d'arbitrage de lignes omnibus, et
 - 40 lesdits moyens d'arbitrage de lignes omnibus recevant des signaux de demande de lignes omnibus individuellement depuis lesdits premiers moyens d'accès à la mémoire, lesdits seconds moyens d'accès à la mémoire et ladite CPU locale et obtenant l'accès auxdits moyens de lignes omnibus locales dans un ordre de priorité prédéterminé.
2. Système de transmission de données numériques selon la revendication 1, caractérisé en outre en ce que lesdits moyens d'arbitrage de lignes omnibus, en réponse à la réception de signaux de demande de lignes omnibus provenant individuellement desdits premiers moyens d'accès à la mémoire, desdits seconds moyens d'accès à la mémoire et de ladite CPU locale, obtiennent l'accès auxdits moyens de lignes omnibus locales dans cet ordre de priorité desdits seconds moyens d'accès à la mémoire ayant la plus haute priorité, lesdits premiers moyens d'accès à la mémoire ayant la priorité supérieure suivante et ladite CPU locale ayant le plus bas ordre de priorité.
3. Système de transmission de données numériques selon la revendication 1 ou 2, caractérisé en outre en ce que lesdits moyens de mémoire locale dudit processeur de transmission comportent une mémoire à lecture/écriture locale (22), lesdits moyens de mémoire locale comportant en outre une mémoire permanente locale (270), et lesdits premiers moyens d'accès à la mémoire connectant lesdits moyens de lignes omnibus de système auxdits moyens de lignes omnibus locales de manière que ladite CPU principale puisse transférer des données vers ladite mémoire à lecture/écriture locale et que ladite CPU locale puisse transférer des données vers ladite mémoire principale.
4. Système de transmission de données numériques selon la revendication 1, caractérisé en outre par des moyens de réception (15) comprenant une entrée (13) couplée avec ledit circuit de signaux et une mémoire tampon de réception (18), ladite mémoire tampon de réception étant couplée avec l'entrée

desdits moyens de réception pour recevoir des données entrantes provenant dudit circuit de signaux, des moyens d'émission (15) comprenant une sortie (14) couplée avec ledit circuit de signaux et une mémoire tampon d'émission (37) destinée à recevoir des données à l'une de ses entrées pour l'émission vers la sortie (14) desdits moyens d'émission, et

5 ledit processeur de transmission comprenant des moyens de commande d'émetteur/récepteur réagissant à la réception de données provenant dudit circuit de signaux dans ledits moyens de réception, ledits seconds moyens d'accès à la mémoire connectant ladite mémoire tampon de réception et ladite mémoire tampon d'émission auxdits moyens de lignes omnibus locales et réagissant auxdits moyens de commande d'émission/réception pour accéder à ladite mémoire à lecture/écriture locale en écrivant 10 des données reçues de la mémoire tampon de réception vers la mémoire à lecture/écriture locale et lisant des données émises par la mémoire à lecture/écriture locale vers ladite mémoire tampon d'émission.

15 5. Système de transmission de données numériques selon la revendication 4, caractérisé en outre en ce qu'une trame de plusieurs multiplets de données à émettre est créée dans ladite mémoire principale à la commande de ladite CPU principale, et cette trame est transférée au déclenchement par ladite CPU principale, par l'intermédiaire desdits moyens de lignes omnibus de système et desdits moyens de lignes omnibus locales, dans ladite mémoire à lecture/écriture locale, et ensuite, cette trame est transférée à la commande desdits moyens de commande d'émission/réception par l'intermédiaire 20 desdits moyens de lignes omnibus locales et de ladite mémoire tampon d'émission vers ledits moyens d'émission.

25 6. Système de transmission de données numériques selon la revendication 5, caractérisé en outre en ce que ledits trames de données créés dans ladite mémoire principale et émises sur ledit circuit de signaux par ledits postes comprennent une adresse de destination identifiant l'un particulier des postes et une adresse de source identifiant le poste émetteur.

30 7. Système de transmission de données numériques selon l'une quelconque des revendications 4 à 6, caractérisé en outre en ce que des données sont émises sur ledit circuit de signaux dans un format de bus en série, ledits moyens de réception comprenant des moyens de conversion série-parallèle (17), ledits moyens d'émission comprenant des moyens de conversion parallèle-série (38) et ledits moyens de lignes omnibus de système et ledits moyens de lignes omnibus locales étant des lignes omnibus en parallèle.

35 8. Système de transmission de données numériques tel ou l'une quelconque des revendications 3 à 7, caractérisé en outre en ce que ladite mémoire permanente locale est couplée avec ledits moyens de lignes omnibus locales, ladite CPU locale exécutant des instructions extraites de ladite mémoire permanente locale par l'intermédiaire desdits moyens de lignes omnibus locales, à la commande desdits moyens d'arbitrage de lignes omnibus et ladite CPU principale exécutant des instructions extraites de ladite mémoire principale par l'intermédiaire desdits moyens de lignes omnibus de système.

40 9. Système de transmission de données numériques tel ou l'une quelconque des revendications 3 à 8, caractérisé en outre en ce que ledits premiers moyens d'accès à la mémoire comprennent des moyens d'entrée/sortie de disposition de mémoire et un registre de commande, ladite CPU principale accédant audit registre de commande par l'intermédiaire desdits moyens d'entrée/ sortie de disposition en mémoire, et
l'accès direct à la mémoire de ladite mémoire à lecture/écriture locale par ledits moyens de lignes omnibus de système étant déclenché par ledits moyens de registre de commande en réponse à 45 l'écriture de données dans ce registre par ladite CPU principale.

50 10. Système de transmission de données numériques selon l'une quelconque des revendications 3 à 9, caractérisé en outre en ce que ledits premiers moyens d'accès à la mémoire comprennent un dispositif d'interface de système (30) qui couple ledits moyens de lignes omnibus locaux avec ledits moyens de lignes omnibus de système pour transférer des données entre ladite mémoire principale et ledits moyens de mémoire locale, et
ledits seconds moyens d'accès à la mémoire comprenant un dispositif d'interface de transmission de données (16) qui couple ledits moyens de lignes omnibus locales avec ledit circuit de signaux pour

transférer des données entre ledit circuit de signaux et leadits moyens de mémoire locale.

11. Système de transmission de données numériques selon la revendication 10, caractérisé en outre en ce que ladite mémoire permanente locale est disposée dans ledit dispositif d'interface de transmission de données et

5 ledit dispositif d'interface de système comprenant un canal d'entrée-sortie de disposition en mémoire par lequel ladite CPU principale accède auxdits moyens de mémoire locale pour transférer une commande et des informations d'état et un canal d'accès direct à la mémoire par lequel des données sont transférées entre la mémoire principale et les moyens de mémoire locale à la commande du dispositif d'interface de système après le déclenchement du transfert par ladite CPU principale ou ladite CPU locale, et

10 ledit dispositif d'interface de transmission de données comprenant un canal d'accès direct à la mémoire par lequel des données sont transférées entre ledit circuit de signaux et ledits moyens de mémoire locale à la commande du dispositif d'interface de transmission de données après le déclenchement du transfert par la CPU locale ou par le dispositif d'interface de transmission de données.

15

12. Système de transmission de données numériques selon la revendication 10 ou 11, caractérisé en outre en ce que ledits moyens de lignes omnibus locales comprennent une ligne omnibus d'adresse/données multiplexée locale (19, 28) qui connecte ladite CPU locale et ledits moyens de mémoire locale, ledits moyens de mémoire locale comprenant ladite mémoire permanente pour mémoriser des instructions de programme exécutées par ladite CPU locale et ladite mémoire à lecture/écriture pour mémoriser momentanément des données de message,

20 ledit dispositif d'interface de système reliant ladite ligne omnibus d'adresse/données locale auxdits moyens de lignes omnibus de système pour transférer des données de message entre ladite mémoire principale et ledits moyens de mémoire locale par accès direct à la mémoire dans les deux directions et

25 ledit dispositif d'interface de transmission de données couplant ledits moyens de lignes omnibus locales avec ledit circuit de signaux pour transférer des données entre ledit circuit de signaux et ledits moyens de mémoire locale par accès direct à la mémoire dans les deux directions.

30

13. Système de transmission de données numériques selon l'une quelconque des revendications précédentes, caractérisé en outre en ce que ledit circuit de signaux est un circuit unilatéral fermé et chaque poste reçoit des données provenant d'une direction en amont dans le circuit de signaux et émettant des données vers ledit circuit de signaux dans une direction en aval,

35 l'accès audit circuit de signaux étant commandé par un jeton inclus dans les données émises et passant d'un poste à un autre autour du circuit fermé.

40

14. Système de transmission de données numériques selon l'une quelconque des revendications précédentes, caractérisé en outre en ce que ledit processeur de transmission consiste en un dispositif à microprocesseur en une seule puce.

Ansprüche

45

1. Digitales Datenübertragungssystem mit:

einem Signalweg (11),

50 mehreren mit dem Signalweg verbundenen Stationen (10), von denen jede digitale Daten von dem Weg empfängt und digitale Daten zu dem Weg sendet, wobei jede Station versehen ist mit

einem Hauptprozessor (12) mit einer Hauptzentraleinheit (31) und einem Hauptspeicher (32) sowie

55 einem Systembusmittel (34), das die Hauptzentraleinheit und den Hauptspeicher miteinander verbindet,

einem Kommunikationsprozessor (20) mit einer Ortszentraleinheit (21) und einem Ortspeichermittel (22, 27), wobei ein Ortsbusmittel (23, 24) vorgesehen ist, das die Ortszentraleinheit und das Ortspeichermittel miteinander verbindet,

5 einem ersten Speicherzugriffsmittel (30), das das Systembusmittel mit dem Ortsbusmittel verbindet, wodurch die Hauptzentraleinheit Daten zu dem Ortspeichermittel übertragen kann und die Ortszentraleinheit Daten zu dem Hauptspeicher übertragen kann, und

10 6 einem zweiten Speicherzugriffsmittel (16), das an das Ortsbusmittel für einen Zugriff auf das Ortsspeichermittel für Datenübertragungsoperationen angeschlossen ist,

15 7 dadurch gekennzeichnet, daß der Kommunikationsprozessor ferner einen Zeitgeber (25) enthält, der ein mit dem Ortsbusmittel verbundenes Zählregistermittel (25d) sowie Mittel (25k, 25p) enthält, die die Ortszentraleinheit unterbrechen, wenn das Zählregistermittel (25d) einen vorbestimmten Wert erreicht, und

20 8 ferner ein Bus-Entscheidungsmittel (26) enthält, das den Zugriff auf das Ortsbusmittel steuert,

25 9 wobei das erste und das zweite Speicherzugriffsmittel wirkungsmäßig mit dem Bus-Entscheidungsmittel verbunden sind, und das Bus-Entscheidungsmittel Busanforderungssignale einzeln von dem ersten Speicherzugriffsmittel, dem zweiten Speicherzugriffsmittel und der Ortszentraleinheit empfängt und den Zugriff auf das Ortsbusmittel in einer vorbestimmten Prioritätsordnung gewährt.

20 10 2. Digitales Datenübertragungssystem nach Anspruch 1, ferner dadurch gekennzeichnet, daß das Bus-Entscheidungsmittel in Abhängigkeit vom Empfang von Busanforderungssignalen von dem ersten Speicherzugriffsmittel, dem zweiten Speicherzugriffsmittel und der Ortszentraleinheit Zugriff auf das Ortsbusmittel in derjenigen Prioritätsordnung gewährt, bei der das zweite Speicherzugriffsmittel die höchste Priorität, das erste Speicherzugriffsmittel die nächsthöhere Priorität und die Orts-Zentraleinheit die niedrigste Priorität hat.

25 11 3. Digitales Datenübertragungssystem nach Anspruch 1 oder 2, ferner dadurch gekennzeichnet, daß das Ortspeichermittel des Kommunikationsprozessors einen Orts-Schreib/Lese-Speicher (22) enthält, daß das Ortspeichermittel ferner einen Ortsfestspeicher(270) enthält und daß das erste Speicherzugriffsmittel das Systembusmittel mit dem Ortsbusmittel verbindet, wodurch die Hauptzentraleinheit Daten zu dem Orts-Schreib/Lese-Speicher übertragen kann und die Ortszentraleinheit Daten zu dem Hauptspeicher übertragen kann.

30 12 4. Digitales Datenübertragungssystem nach Anspruch 1, ferner gekennzeichnet durch ein Empfangsmittel (15) mit einem Eingang (13), der mit dem Signalweg und einem Empfangspufferspeicher (18) verbunden ist, wobei der Empfangspufferspeicher mit dem Eingang des Empfangsmittels zum Empfangen ankommender Daten von dem Signalweg verbunden ist, ein Sendemittel (15) mit einem Ausgang (14), der mit dem Signalweg verbunden ist, und einen Sendepufferspeicher (37) zum Empfangen von Daten an einem Eingang für das Aussenden zu dem Ausgang (14) des Sendemittels, wobei der Kommunikationsprozessor Send/Empfangs-Steuermittel enthält, die abhängig vom Empfang von Daten von dem Signalweg in dem Empfangsmittel arbeiten, wobei das zweite Speicherzugriffsmittel den Empfangspufferspeicher und den Sendepufferspeicher mit dem Ortsbusmittel verbindet und abhängig von den Sender/Empfangs-Steuermitteln einen Zugriff auf den Orts-Schreib/Lese-Speicher zum Schreiben empfangener Daten von dem Empfangspufferspeicher in den Orts-Schreib/Lese-Speicher und zum Lesen von Sendedaten von dem Orts-Schreib/Lese-Speicher zu dem Sendepufferspeicher ausüben.

35 13 5. Digitales Datenübertragungssystem nach Anspruch 4, ferner dadurch gekennzeichnet, daß in den Hauptspeicher unter der Steuerung durch die Hauptzentraleinheit ein Rahmen aus mehreren zu sendenden Daten-Bytes gebildet wird und daß dieser Rahmen nach Auslösung durch die Hauptzentraleinheit über das Systembusmittel und das Ortsbusmittel in den Orts-Schreib/ Lese-Speicher übertragen werden, wobei der Rahmen dann oder der Steuerung durch die Send/Empfangs-Steuermittel über das Ortsbusmittel und den Sendepufferspeicher zu dem Sendemittel übertragen wird.

40 14 6. Digitales Datenübertragungssystem nach Anspruch 5, ferner dadurch gekennzeichnet, daß die in dem Hauptspeicher erzeugten und auf dem Signalweg durch die Stationen übertragenen Datenrahmen eine Bestimmungsadresse enthalten, die eine bestimmte Station identifiziert, sowie eine Quellenadresse enthalten, die die Ursprungsstation identifiziert.

7. Digitales Datenübertragungssystem nach einem der Ansprüche 4 bis 6, ferner dadurch gekennzeichnet, daß Daten auf dem Signalweg in einem Bit-serialen Format übertragen werden, daß das Empfangsmit-
tel ein Serien-Parallel-Umsetzungsmittel (17) enthält, daß das Sondermittel ein Parallel-Serien-Umset-
zungsmittel (38) enthält und daß das Systembusmittel und das Ortsbusmittel parallele Busmittel sind.
8. Digitales Datenübertragungssystem nach einem der Ansprüche 3 bis 7, ferner dadurch gekennzeichnet, daß der Orts-Schreib/Lese-Speicher mit dem Ortsbusmittel verbunden ist, wobei die Ortszentraleinheit aus dem Orts-Schreib/Lese-Speicher über das Ortsbusmittel unter der Steuerung durch das Entschei-
dungsmittel abgerufene Befehle ausführt, und daß die Hauptzentraleinheit aus dem Hauptspeicher über
das Systembusmittel abgerufene Befehle ausführt.
9. Digitales Datenübertragungssystem nach einem der Ansprüche 3 bis 8, ferner dadurch gekennzeichnet, daß das ersten Speicherzugriffsmittel speicherorientierte Eingabe/Ausgabe-Mittel und ein Steuerregister enthält, wobei die Hauptzentraleinheit einen Zugriff auf das Steuerregister über die speicherorientierten
Eingabe/Ausgabe-Mittel ausübt, und daß ein direkter Speicherzugriff auf den Orts-Schreib/Lese-Spei-
cher über das Systembusmittel von dem Steuerregister in Abhängigkeit vom Schreiben von Daten in
dieses Register durch die Hauptzentraleinheit ausgelöst wird.
10. Digitales Datenübertragungssystem nach einem der Ansprüche 3 bis 9, ferner dadurch gekennzeichnet,
daß das erste Speicherzugriffsmittel eine System-Schnittstellenvorrichtung (30) enthält, die das Ortsbus-
mittel mit dem Systembusmittel zum Übertragen von Daten zwischen dem Hauptspeicher und dem
Ortspeichermittel koppelt, und daß das zweite Speicherzugriffsmittel eine Datenübertragungs-Schnit-
stellenvorrichtung (16) enthält, die das Ortsbusmittel mit dem Signalweg für die Datenübertragung
zwischen dem Signalweg und dem Ortspeichermittel koppelt.
11. Digitales Datenübertragungssystem nach Anspruch 10, ferner dadurch gekennzeichnet, daß der Orts-
Schreib/Lese-Speicher in der Datenübertragungs-Schnittstellenvorrichtung untergebracht ist, daß die System-Schnittstellenvorrichtung einen speicherorientierten Eingabe/Ausgabe-Kanal enthält, durch den die Hauptzentraleinheit einen Zugriff auf das Ortspeichermittel zum Übertragen von Befehls- und
Statusinformationen ausübt, und ferner einen Direkt-Speicherzugriffskanal enthält, durch den Daten
zwischen dem Hauptspeicher und dem Ortspeichermittel unter der Steuerung durch die System-
schnittstellenvorrichtung nach Auslösung der Übertragung durch die Hauptzentraleinheit oder die
Ortszentraleinheit übertragen werden, und daß die Datenübertragungs-Schnittstellenvorrichtung einen
Direkt-Speicherzugriffskanal enthält, durch den Daten zwischen dem Signalweg und dem Ortsspeicher-
mittel unter der Steuerung durch die Datenübertragungs-Schnittstellenvorrichtung nach der Auslösung
der Übertragung durch die Ortszentraleinheit oder die Datenübertragungs-Schnittstellenvorrichtung
übertragen werden.
12. Digitales Datenübertragungssystem nach Anspruch 10 oder 11, ferner dadurch gekennzeichnet, daß
das Ortsbusmittel einen örtlichen, multiplexierten Adressen/Daten-Bus (19, 28) enthält, der die Ortszen-
traleinheit mit dem Ortspeichermittel verbindet, wobei das Ortspeichermittel den Schreib/Lese-
Speicher zum Speichern von durch die Ortszentraleinheit ausgeführten Programmbefehlen und den
Schreib/Lese-Speicher zum Zwischenspeichern von Nachrichtendaten enthält, daß die System-Schnit-
stellenvorrichtungen örtlichen Adressen/Daten-Bus mit den Systembusmitteln zum Übertragen von
Nachrichtendaten zwischen dem Hauptspeicher und dem Ortspeichermittel durch direkten Speicher-
zugriff in beiden Richtungen koppelt, und daß die Datenübertragungs-Schnittstellenvorrichtung das
Ortsbusmittel zum Übertragen von Daten zwischen dem Signalweg und dem Ortspeichermittel durch
direkten Speicherzugriff in beiden Richtungen mit dem Signalweg koppelt.
13. Digitales Datenübertragungssystem nach einem der vorhergehenden Ansprüche, ferner dadurch ge-
kennzeichnet, daß der Signalweg ein geschlossener, in einer Richtung wirksamer Weg ist, daß jede
Station Daten aus einer Zuleitungsrichtung im Signalweg empfängt und Daten in einer Ableitungsrich-
tung an den Signalweg abgibt und daß der Zugriff auf den Signalweg durch ein in den übertragenen
Daten enthaltenes Kennzeichen gesteuert wird, das von Station zu Station durch den geschlossenen
Weg geschickt wird.
14. Digitales Datenübertragungssystem nach einem der vorhergehenden Ansprüche, ferner dadurch ge-
kennzeichnet, daß der Kommunikationsprozessor eine Ein-Chip-Mikroprozessorenvorrichtung ist.




Fig. 1




Fig. 1a




Fig. 1b




Fig. 2




Fig.2a




Fig.2b

Fig.2c




Fig.2d





Fig.3







Fig.3a



CHIP 20 — Fig. 4






Fig. 4a








Fig. 5c




Fig. 5d




Fig. 5e




Fig. 5f




Fig. 5g




Fig. 6

STATE TIME	OPERATION	ALU INPUTS/OUTPUT	ADDRESS/DATA BUS 4	OTHER
S1	PRE-FETCH INSTRUCTION AT ADDRESS IN PC FOR THE ADD OPERATION; PERFORM ALU OPERATION FOR OLD INSTRUCTION	PERFORMS ALU OPERATION FOR OLD INSTRUCTION. RESULT TO E BUS DELAYED	READ CYCLE: ADDRESS FROM PC OUT VIA P BUS AND BUS 4; DATA BACK IN ON BUS 4 TO K REG VIA BUS 5	INCREMENT PC DEN-ACTIVE
S2	WRITE RESULT OF OLD INSTRUCTION AT OLD DESTINATION ADDRESS. CALCULATE SOURCE ADDRESS FOR NEW INSTRUCTION	ADDS 25 + WP; A INPUT IS D FIELD FROM PREVIOUS INSTRUCTION WORD OFF DI BUS; B INPUT IS WP VIA S BUS. DUTPUT TO E, DELAYED	WRITE CYCLE: DEST ADDRESS OUT FROM D REG. INSTRUCTION OUT ON BUS 4 FROM MA. THEN DATA OUT ON BUS 4 FROM D REG.	LOAD IR VIA DI BUS AND DETERMINE ENTRY POINT FOR ADD INSTRUCTION
S3	FETCH SOURCE OPERAND AT ADDRESS IN MA; CALCULATE ADDRESS OF DESTINATION REGISTER	ADDS 2D + WP; A INPUT IS D FIELD FROM INSTRUCTION; B INPUT IS WP; RESULT TO MA VIA E DELAYED	READ CYCLE: SOURCE ADDRESS OUT FROM MA TO BUS 4; DATA BACK IN ON BUS 4 TO K REG	DEN-ACTIVE
S4	FETCH CONTENTS OF DESTINATION REGISTER	ADDS 2D + WP AGAIN; RESULT TO MA TO BE USED AS THE DESTINATION WRITE	READ CYCLE: ADDRESS (2D + WP) GOES OUT CONTENTS OF THAT ADDRESS BACK TO BE USED FOR NEXT FETCH	DEN-ACTIVE
S5	FETCH DESTINATION OPERAND	ADDS K PLUS ZERO AND STORES RESULT IN T REG	READ CYCLE: DESTINATION ADDRESS FROM MA GOES OUT ON BUS 4; DEST OPERAND COMES BACK ON BUS 4 INTO K REG	DEN-ACTIVE
S6	PERFORM THE ADD OPERATION FOR THIS INSTRUCTION, AND PRE-FETCH THE CONTENTS OF ADDRESS IN PC FOR THE NEXT INSTRUCTION	PERFORMS ADD OPERATION FOR CURRENT ADD INSTRUCTION. INPUT FROM T, B INPUT FROM T VIA F, OUTPUT TO E DELAYED	READ CYCLE: ADDRESS FOR NEXT INSTRUCTION OUT VIA P BUS AND BUS 4; DATA BACK IN ON BUS 4 TO K REG VIA BUS 5	INCREMENT PC DEN-ACTIVE
S7	WRITE RESULT OF THIS ADD OPERATION AT DESTINATION ADDRESS. CALCULATE SOURCE ADDRESS FOR NEXT INSTRUCTION	ADDS 25 + WP FOR NEXT INSTRUCTION	WRITE CYCLE: DESTINATION ADDRESS FOR ADD INSTRUCTION GOES OUT ON BUS 4 FROM MA. THEN DATA ON TO BUS 4 FROM D REG.	LOAD IR VIA DI BUS AND DETERMINE ENTRY POINT FOR NEXT INSTRUCTION

Fig. 7




Fig. 7a




Fig. 8




Fig. 9